

A CORRESPONDENCE PRINCIPLE BETWEEN (HYPER)GRAPH THEORY AND PROBABILITY THEORY, AND THE (HYPER)GRAPH REMOVAL LEMMA

TERENCE TAO

ABSTRACT. We introduce a correspondence principle (analogous to the Furstenberg correspondence principle) that allows one to extract an infinite random graph or hypergraph from a sequence of increasingly large deterministic graphs or hypergraphs. As an application we present a new (infinitary) proof of the hypergraph removal lemma of Nagle-Schacht-Rödl-Skokan and Gowers, which does not require the hypergraph regularity lemma and requires significantly less computation. This in turn gives new proofs of several corollaries of the hypergraph removal lemma, such as Szemerédi's theorem on arithmetic progressions.

1. INTRODUCTION

It is an interesting phenomenon in mathematics that certain types of problems can be treated both by finitary means (e.g. using combinatorial analysis of finite sets), and by infinitary means (e.g. using constructions involving the axiom of choice), thus giving parallel but distinct ways to prove a single result. One particularly striking example of this is *Szemerédi's theorem* (see Theorem 2.1) on arithmetic progressions. This difficult and important theorem now has several proofs, both finitary and infinitary, using fields of mathematics as diverse as Fourier analysis, ergodic theory, graph theory, hypergraph theory, and elementary combinatorics; the finitary and infinitary arguments are connected by the beautiful *Furstenberg correspondence principle* (see Section 2). These proofs have different strengths and weaknesses; generally speaking, the infinitary proofs are cleaner, shorter, and more elegant, but require significantly more machinery, whereas the finitary proofs are more elementary and provide more quantitative results, but tend to be messier and longer in nature. One particularly visible difference is that finitary proofs often require a number of small parameters (such as ε, δ) or large parameters (such as N, M), whereas in the infinitary analogues of these proofs, the small parameters often have become zero and the large parameters have become infinite, which can lead to cleaner (but more subtle) arguments.

Some progress has been made in reconciling the finitary and infinitary approaches¹, as it has been increasingly realized that ideas and methods from the infinitary world

¹From a proof-theoretical perspective, one can use quantifier-elimination methods (such as Herbrand's theorem) to automatically convert a large class of infinitary arguments to finitary ones; this was for instance carried out for the Furstenberg-Weiss infinitary proof of van der Waerden's

can be transferred to the finitary world, and vice versa; see for instance [29] for a finitary version of the infinitary ergodic approach to Szemerédi’s theorem. Such a fusion of ideas from both sources proved to be particularly crucial in the recent result [13] that the primes contained arbitrarily long progressions; this argument was almost entirely finitary in nature, yet at the same time it relied heavily on ideas from the infinitary world of ergodic theory (see [17], [15] for further discussion of this connection).

In this paper we investigate a transference in the other direction, taking results from the finitary world of combinatorics (and in particular graph theory and hypergraph theory), and identifying them with a corresponding result in the infinitary world, which in this case turns out to be the world of probability theory² (or measure theory). In particular, we present a *correspondence principle*, analogous to the Furstenberg correspondence principle, that shows how any sequence of increasingly large graphs or hypergraphs has a “weak limit”, which we view as an infinitely large *random* graph or hypergraph³. This principle is slightly more complicated than the Furstenberg correspondence principle, but does not use the full power of deep results such as the Szemerédi regularity lemma or its extension to hypergraphs; indeed we do not explicitly state or use such a regularity lemma in this work here, although ideas from that lemma are certainly involved in several components of the argument.

The main advantage of passing from a deterministic finite graph to a random infinite graph is that one now obtains a number of *factors* (σ -algebras) in the probability space which enjoy some very useful invariance and relative independence properties. One can think of the presence of these factors as being analogous to the partitions obtained by the Szemerédi regularity lemma that make a graph ε -regular, but with the distinction that the partition is now infinite and the ε parameter set to zero (so one now has perfect regularity). This sending of the epsilon parameters to zero turns out to be extremely useful in cleaning up proofs of certain statements which previously could only be proven via a regularity lemma. In particular, we will give an infinitary proof here of the *triangle removal lemma* of Ruzsa and Szemerédi [24], as well as the substantially more difficult *hypergraph removal lemma* of Nagle, Rödl, Schacht, and Skokan [19], [20], [22], [23] and Gowers [12] (as well as a later refinement in [30]). As this lemma is already strong enough to deduce Szemerédi’s theorem on arithmetic progressions (as well as a multidimensional generalisation due to Furstenberg and Katznelson [9]), we have thus presented yet another proof of Szemerédi’s theorem here. These lemmas have some further applications; for

theorem via topological dynamics, see [11]. However such methods do not seem to shed much light on the connection between the infinitary proofs and the existing finitary proofs in the literature.

²This is actually not all that surprising, given that *finitary* probability theory has already proven to have a major role to play in graph theory.

³This is related to, but slightly different from, a different concept of graph limit developed by Lovász and Szegedy in [18], in which the limiting object becomes a “continuous weighted graph”, or more precisely a symmetric measurable function from $[0, 1] \times [0, 1]$ to $[0, 1]$. Such a concrete limiting object is particularly useful for computations such as counting the number of induced subgraphs of a certain shape; it also can be used to establish results such as the triangle removal lemma (Szegedy, personal communication).

instance, they were used in [31] to show that the Gaussian primes contained arbitrarily shaped constellations. In Appendix B we discuss the connections (or lack thereof) between these infinitary removal lemmas, and the recurrence theorems of Furstenberg and later authors.

The setting of this paper was deliberately placed at a midpoint between graph theory and ergodic theory, and the author hopes that it illuminates the analogies and interconnections between these two subjects.

The author thanks Balász Szegedy for many useful discussions, Timothy Gowers for suggesting the original topic of investigation, Vitaly Bergelson for encouragement, and Olivier Gerard for corrections. The author is especially indebted to the anonymous referees for many corrections and suggestions. The author is supported by a grant from the Packard Foundation.

2. MOTIVATION: THE FURSTENBERG CORRESPONDENCE PRINCIPLE

To motivate the correspondence principle for graphs and hypergraphs, we first review the Furstenberg correspondence principle which connects results such as Szemerédi's theorem with recurrence results in ergodic theory. Let us recall Szemerédi's theorem in a quantitative (finitary) form:

Theorem 2.1 (Szemerédi's theorem, quantitative version). [26] *Let $0 < \delta \leq 1$ and $k \geq 1$. Let A be a subset of a cyclic group $\mathbf{Z}_N := \mathbf{Z}/N\mathbf{Z}$ whose cardinality $|A|$ is at least δN . Then there exist at least $c(k, \delta)N^2$ pairs $(x, r) \in \mathbf{Z}_N \times \mathbf{Z}_N$ such that $x, x + r, \dots, x + (k - 1)r \in A$, where $c(k, \delta) > 0$ is a positive quantity depending only on k and δ .*

This result is easily seen to imply to Szemerédi's theorem in its traditional (infinitary) form, which asserts that every set of integers of positive upper density contains arbitrarily long progressions. The converse implication also follows from an argument of Varnavides [33]. This particular formulation of Szemerédi's theorem played an important role in the recent result [13] that the primes contained arbitrarily long arithmetic progressions.

In 1977, Furstenberg obtained a new proof of Szemerédi's theorem by deducing it from the following result in ergodic theory.

Theorem 2.2 (Furstenberg recurrence theorem). [7], [10] *Let $(\Omega, \mathcal{B}_{\max}, \mathbf{P})$ be a probability space (see Appendix A for probabilistic notation). Let $T : \Omega \rightarrow \Omega$ be a bi-measurable map which is probability preserving, thus $\mathbf{P}(T^n A) = \mathbf{P}(A)$ for all events $A \in \mathcal{B}_{\max}$ and $n \in \mathbf{Z}$. Then for all $k \geq 1$ and all events $A \in \mathcal{B}_{\max}$ with $\mathbf{P}(A) > 0$, we have*

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \mathbf{P}(A \wedge T^n A \wedge \dots \wedge T^{(k-1)n} A) > 0.$$

The deduction of Theorem 2.1 from Theorem 2.2 proceeds by the *Furstenberg correspondence principle* [7], [10], [8]. Let us give a slightly non-standard exposition of this principle (in particular drawing heavily on the language of probability theory), in order to motivate an analogous principle for graphs and hypergraphs in later sections. We shall interpret this correspondence principle as an assertion that any sequence $(A^{(m)}, \mathbf{Z}_{N^{(m)}})$ of sets $A^{(m)}$ in a cyclic group $\mathbf{Z}_{N^{(m)}}$ can have an asymptotic limit as $m \rightarrow \infty$, which will end up being a probability space endowed with a probability-preserving shift T . To state this more precisely we shall need some notation. First, we describe a certain universal space in which it will be convenient to take limits.

Definition 2.3 (Furstenberg universal space). Let $\Omega := 2^{\mathbf{Z}} := \{B : B \subset \mathbf{Z}\}$ denote the set of all subsets $B \subset \mathbf{Z}$ of the integers \mathbf{Z} ; one can also view this space as the infinite cube $\{0, 1\}^{\mathbf{Z}}$ if desired. We give this space the product σ -algebra \mathcal{B}_{\max} , generated by the events⁴ $A_n := \{B \in \Omega : n \in B\}$ for $n \in \mathbf{Z}$. Indeed, one can think of $(\Omega, \mathcal{B}_{\max})$ as being the universal event space generated by the countable sequence of events A_n . The space Ω enjoys an obvious shift action $T : \Omega \rightarrow \Omega$, defined by $TB := B + 1 := \{n + 1 : n \in B\}$ for all $B \in \Omega$. This then induces a shift $T : \mathcal{B}_{\max} \rightarrow \mathcal{B}_{\max}$ in the obvious manner, thus for instance $T^n A_0 = A_n$. We define the *regular algebra* \mathcal{B}_{reg} of \mathcal{B}_{\max} to be the algebra generated by the A_n , thus the events in \mathcal{B}_{reg} (which we refer to as *regular events*) are those events which are generated by at most finitely many of the A_n (i.e. those events that only require knowing the truth value of $n \in B$ for finitely many values of n).

Now we embed finite objects $(A^{(m)}, \mathbf{Z}_{N^{(m)}})$ described earlier in this universal space.

Definition 2.4 (Furstenberg universal embedding). Let $m \geq 1$, let $\mathbf{Z}_{N^{(m)}}$ be a cyclic group with $N^{(m)} \geq m$, and let $A^{(m)}$ be a subset of $\mathbf{Z}_{N^{(m)}}$. We define the probability space $(\Omega^{(m)}, \mathcal{B}_{\max}^{(m)}, \mathbf{P}^{(m)})$ as the space corresponding to sampling⁵ $x^{(m)}$ and $\lambda^{(m)}$ uniformly and independently at random from $\mathbf{Z}_{N^{(m)}}$ and $[L^{(m)}]$, where $[N] := \{1, \dots, N\}$ denotes the integers from 1 to N and $L^{(m)} \geq 1$ is the integer part of $N^{(m)}/m$. We then map every pair $(x^{(m)}, \lambda^{(m)})$ of $\Omega^{(m)}$ to a point $B^{(m)} \in \Omega$ (i.e. a subset of the integers \mathbf{Z}) by the formula

$$B^{(m)} := \{n \in \mathbf{Z} : x^{(m)} + n\lambda^{(m)} \in A^{(m)}\};$$

one can think of this as a random lifting of the set $A^{(m)} \subset \mathbf{Z}_{N^{(m)}}$ up to the integers \mathbf{Z} . This mapping from $\Omega^{(m)}$ to Ω is clearly measurable, since the inverse images of the generating events A_n in $(\Omega, \mathcal{B}_{\max})$ are simply the events that $x^{(m)} + n\lambda^{(m)} \in A^{(m)}$, which are certainly measurable in $\mathcal{B}_{\max}^{(m)}$. This allows us to extend the probability measure $\mathbf{P}^{(m)}$ from $(\mathcal{B}_{\max}^{(m)}, \Omega^{(m)})$ to the product space $(\mathcal{B}_{\max} \times$

⁴A more topological of thinking about this proceeds by endowing Ω with the product topology, so that it becomes a compact Hausdorff totally disconnected space, and then letting \mathcal{B}_{\max} be the Borel σ -algebra, generated by the open sets. The regular algebra \mathcal{B}_{reg} then consists of those events which are simultaneously open and closed, or equivalently those events whose indicator function is continuous.

⁵The introduction of the dilation parameter $\lambda^{(m)}$ is essentially the averaging trick of Varnavides [33]. The exact construction of this space is not important so long as one has the independent random variables $x^{(m)}$ and $\lambda^{(m)}$, but for sake of concreteness one can set $\Omega^{(m)} := \mathbf{Z}_{N^{(m)}} \times [L^{(m)}]$, $\mathcal{B}_{\max}^{(m)} := 2^{\Omega^{(m)}}$ to be the power set of $\Omega^{(m)}$, and $\mathbf{P}^{(m)}$ to be the uniform distribution on $\Omega^{(m)}$.

$\mathcal{B}_{\max}^{(m)}, \Omega \times \Omega^{(m)}$ in a canonical manner⁶, identifying the events A_n with the events $x^{(m)} + n\lambda^{(m)} \in A^{(m)}$. We shall abuse notation and refer to the extended measure also as $\mathbf{P}^{(m)}$.

In more informal terms, the Furstenberg embedding has created, for each m , a random set $B^{(m)} \subset \mathbf{Z}$ which will capture all the important information about the original set $A^{(m)}$ and $\mathbf{Z}_{N^{(m)}}$. For instance, the density of $A^{(m)}$ is nothing more than the probability that 0 lies in $B^{(m)}$, or equivalently the probability of the event A_0 . One can view $x^{(m)}$ and $\lambda^{(m)}$ as the “hidden variables” which generate this random set $B^{(m)}$. However, in order to invoke the correspondence principle we will need to “forget” that the random set $B^{(m)}$ actually came from these variables; indeed, we are going to restrict $\mathbf{P}^{(m)}$ to the common factor $(\Omega, \mathcal{B}_{\max})$ in order to take limits as $m \rightarrow \infty$. More precisely, we have

Proposition 2.5 (Furstenberg correspondence principle). *For every $m \geq 1$, let $\mathbf{Z}_{N^{(m)}}$ be a cyclic group with $N^{(m)} \geq m$, and let $A^{(m)}$ be a subset of $\mathbf{Z}_{N^{(m)}}$, and let $\mathbf{P}^{(m)}$ be as in Definition 2.4. Then there exists a subsequence $0 < m_1 < m_2 < \dots$ of m , and a probability measure $\mathbf{P}^{(\infty)}$ on the Furstenberg universal space $(\Omega, \mathcal{B}_{\max})$, such that we have the weak convergence property*

$$\lim_{i \rightarrow \infty} \mathbf{P}^{(m_i)}(E) = \mathbf{P}^{(\infty)}(E) \text{ for all } E \in \mathcal{B}_{\text{reg}}. \quad (1)$$

Furthermore, we have the shift invariance property

$$\mathbf{P}^{(\infty)}(T^n E) = \mathbf{P}^{(\infty)}(E) \text{ for all } E \in \mathcal{B}_{\max}, n \in \mathbf{Z}. \quad (2)$$

Proof The algebra \mathcal{B}_{reg} is countable. Thus the existence of the weak limit $\mathbf{P}^{(m_i)}$ follows from Lemma A.15. Now observe that the random sets $B^{(m)}$ and $T^n B^{(m)} = B^{(m)} + n$ have the same probability distribution (because x_0 and $x_0 - n\lambda$ have the same distribution for any fixed λ). Thus we observe that $\mathbf{P}^{(m)}$ is shift-invariant:

$$\mathbf{P}^{(m)}(T^n E) = \mathbf{P}^{(m)}(E) \text{ for all } E \in \mathcal{B}_{\max}, n \in \mathbf{Z}.$$

Applying (1) we obtain (2) for all *regular* events E . But since $\mathbf{P}^{(\infty)}$ is countably additive, we see that the space of events E for which (2) holds for every T is a σ -algebra which contains \mathcal{B}_{reg} , and thus contains \mathcal{B}_{\max} as claimed. \blacksquare

Now we can deduce Theorem 2.1 from Theorem 2.2.

Proof [of Theorem 2.1 assuming Theorem 2.2] Suppose that Theorem 2.1 fails. Then we can find $k \geq 1$ and $0 < \delta \leq 1$, a sequence $N^{(m)}$ of positive integers, and a sequence of sets $A^{(m)} \subset \mathbf{Z}_{N^{(m)}}$ of density $|A^{(m)}|/N^{(m)} \geq \delta$ such that

$$\lim_{m \rightarrow \infty} \frac{1}{(N^{(m)})^2} |\{(x, r) \in \mathbf{Z}_{N^{(m)}} : x, x+r, \dots, x+(k-1)r \in A^{(m)}\}| = 0.$$

⁶More precisely, we graph the measurable mapping from $\Omega^{(m)}$ to Ω as a measurable mapping from $\Omega^{(m)}$ to $\Omega \times \Omega^{(m)}$, which contravariantly induces a σ -algebra homomorphism from the product σ -algebra $\mathcal{B}_{\max} \times \mathcal{B}_{\max}^{(m)}$ to $\mathcal{B}_{\max}^{(m)}$. Pulling back the probability measure $\mathbf{P}^{(m)}$ under this homomorphism yields the extension. A similar construction applies to the graph and hypergraph embeddings in later sections.

By passing to a subsequence of m if desired we can make this convergence arbitrarily fast; for instance, we can ensure that

$$\frac{1}{(N^{(m)})^2} |\{(x, r) \in \mathbf{Z}_{N^{(m)}} : x, x+r, \dots, x+(k-1)r \in A^{(m)}\}| \leq \delta 100^{-m}. \quad (3)$$

Observe that the left-hand side is at least $\frac{1}{(N^{(m)})^2} |A^{(m)}| \geq \delta/N^{(m)}$, so we conclude that

$$N^{(m)} \geq 100^m.$$

In particular $N^{(m)} \geq m$, so we can invoke the Furstenberg correspondence principle and obtain a shift-invariant system $(\Omega, \mathcal{B}_{\max}, \mathbf{P}^{(\infty)})$ on the Furstenberg universal space $(\Omega, \mathcal{B}_{\max})$ with the stated properties.

Now let us compute some probabilities in this system, starting with the probability of A_0 . From definition of $\mathbf{P}^{(m)}$ we have

$$\mathbf{P}^{(m)}(A_0) = \mathbf{P}^{(m)}(x_0 \in A^{(m)}) = |A^{(m)}|/N^{(m)} \geq \delta$$

so by (1) we have

$$\mathbf{P}(A_0) \geq \delta.$$

In particular A_0 has strictly positive probability. Next, let $1 \leq n \leq m$ and consider the expression

$$\begin{aligned} \mathbf{P}^{(m)}(A_0 \wedge T^n A_0 \wedge \dots \wedge T^{(k-1)n} A_0) &= \mathbf{P}^{(m)}(A_0 \wedge A_n \wedge \dots \wedge A_{(k-1)n}) \\ &= \mathbf{P}^{(m)}(x^{(m)} + jn\lambda^{(m)} \in A^{(m)} \ \forall 0 \leq j < k) \\ &= \frac{1}{N^{(m)}L^{(m)}} |\{(x^{(m)}, \lambda^{(m)}) \in \mathbf{Z}_{N^{(m)}} \times [L^{(m)}] : \\ &\quad x^{(m)} + jn\lambda^{(m)} \in A^{(m)} \ \forall 0 \leq j < k\}|. \end{aligned}$$

Now observe from definition of $L^{(m)}$, the progressions $x^{(m)}, x^{(m)} + n\lambda^{(m)}, \dots, x^{(m)} + (k-1)n\lambda^{(m)}$ are all distinct as $x^{(m)}$ and $\lambda^{(m)}$ vary. Applying (3), we see that

$$\begin{aligned} \mathbf{P}^{(m)}(A_0 \wedge T^n A_0 \wedge \dots \wedge T^{(k-1)n} A_0) &= \mathbf{P}^{(m)}(A_0 \wedge A_n \wedge \dots \wedge A_{(k-1)n}) \\ &\leq \frac{1}{N^{(m)}L^{(m)}} (N^{(m)})^2 \delta 100^{-m} \\ &\leq 2m 100^{-m} \end{aligned}$$

(say) for all $1 \leq n \leq m$. In particular we have

$$\lim_{m \rightarrow \infty} \mathbf{P}^{(m)}(A_0 \wedge T^n A_0 \wedge \dots \wedge T^{(k-1)n} A_0) = 0$$

for each fixed $n \geq 1$, and hence by (1)

$$\mathbf{P}^{(\infty)}(A_0 \wedge T^n A_0 \wedge \dots \wedge T^{(k-1)n} A_0) = 0$$

for all $n \geq 1$. But this contradicts Theorem 2.2. This completes the deduction of Theorem 2.1 from Theorem 2.2. \blacksquare

Remark 2.6. Note that as this proof proceeded by contradiction, it does not obviously give any sort of *quantitative* lower bound for the quantity $c(k, \delta)$ appearing in Theorem 2.1. It is actually possible (with nontrivial effort) to extract such a bound by taking the *proof* of Theorem 2.2 and making everything finitary; see [29].

However the bounds obtained in this manner are extremely poor. The same remarks apply to the infinitary proofs of the triangle removal lemma and hypergraph removal lemma that we give below. As a related remark, observe that the above argument, while infinitary, did not require the axiom of choice, as one can eliminate the apparent use of choice at the beginning of the argument by well-ordering the objects A , \mathbf{Z}_N , δ appearing in Theorem 2.1 in some standard manner. (The use of Lemma A.15 also does not require choice; see Remark A.16. The original proof of the Furstenberg recurrence theorem in [7] is also choice-free, though the later proof in [10] is not, as it uses Zorn's lemma.) Indeed we will not actually need the axiom of choice in this entire paper, though we shall assume it in order to simplify the exposition slightly.

Remark 2.7. One can also reverse the above argument, and use Theorem 2.1 to deduce Theorem 2.2, basically by applying Theorem 2.1 to various truncated versions of the random set $B := \{n \in \mathbf{Z} : T^n x \in A\}$, where x is sampled from the sample space Ω using the probability measure \mathbf{P} . We omit the standard details.

3. THE GRAPH CORRESPONDENCE PRINCIPLE

We now develop an analogue of the Furstenberg correspondence principle for graphs; namely, we start with a sequence of (undirected) graphs $G^{(m)} = (V^{(m)}, E^{(m)})$ for each $m \geq 1$, and wish to extract (after passing to a subsequence of m 's) some sort of infinitary weak limit. This type of problem was already addressed in [18], with the main tool being a certain weak form of the Szemerédi regularity lemma. Our approach is somewhat similar (though not identical), and the regularity lemma will appear only after the infinite limit is extracted, in Lemma 3.5 below.

As before, we need a universal space in which to take limits. Just as the Furstenberg universal space consisted of infinite *sets* of integers, the graph universal space will consist of infinite *graphs* on the natural numbers. The shift T (which represents a \mathbf{Z} -action) is now replaced⁷ by the action of the permutation group S_∞ , defined as the group of all permutations $\sigma : \mathbf{Z} \rightarrow \mathbf{Z}$ of the integers.

Definition 3.1 (Graph universal space). Let $\mathbf{N} := \{1, 2, \dots\}$ denote the natural numbers, and let $\Omega := 2^{\binom{\mathbf{N}}{2}} = \{(\mathbf{N}, E_\infty) : E_\infty \subset \binom{\mathbf{N}}{2}\}$ denote the space of all (infinite) graphs (\mathbf{N}, E_∞) on the natural numbers, thus the edge set E_∞ is an arbitrary collection of unordered pairs of distinct integers. On this space Ω , we introduce the events $A_{i,j} = A_{j,i}$ for any unordered pair of distinct natural numbers $\{i, j\} \in \binom{\mathbf{N}}{2}$ by $A_{i,j} := \{(\mathbf{N}, E_\infty) \in \Omega : (i, j) \in E_\infty\}$, and let \mathcal{B}_{\max} be the σ -algebra generated by the countable sequence of events $A_{i,j}$. (We adopt the convention that $A_{i,i} = \emptyset$ for all $i \in \mathbf{N}$, thus our graphs have no loops.) We also introduce the regular algebra \mathcal{B}_{reg} generated by the $A_{i,j}$, thus these are the events that depend only on finitely many of the $A_{i,j}$. For any permutation $\sigma : \mathbf{N} \rightarrow \mathbf{N}$ of the natural numbers, we define the associated action on \mathcal{B}_{\max} by mapping $\sigma : A_{i,j} \mapsto A_{\sigma(i), \sigma(j)}$ and extending this to a σ -algebra isomorphism in the unique manner. More explicitly, σ will map each graph (N, E_∞) to the graph $(N, \sigma E_\infty)$, where $\sigma E_\infty := \{\{\sigma(i), \sigma(j)\} :$

⁷We are indebted to Balász Szegedy for pointing out the analogy between the \mathbf{Z} -action of a dynamical system and the S_∞ -action on an infinite graph.

$\{i, j\} \in E_\infty\}$. For any (possibly infinite) subset I of \mathbf{N} , we define \mathcal{B}_I to be the factor of \mathcal{B}_{\max} generated by the events $A_{i,j}$ for $i, j \in I$; informally speaking, \mathcal{B}_I represents the knowledge obtained by measuring the restriction of E_∞ to I . Observe the trivial monotonicity $\mathcal{B}_I \subseteq \mathcal{B}_J$ whenever $I \subseteq J$.

The space $(\Omega, \mathcal{B}_{\max})$ is thus the universal event space associated to the events $A_{i,j}$, and is the natural event space for studying infinite random graphs. (For instance, the infinite Erdős-Renyi random graph $G(\infty, p)$ for fixed $0 \leq p \leq 1$, where the vertex set is \mathbf{Z} (say) and any two integers are connected by an edge with an independent probability of p , would correspond to the scenario in which all the events $A_{i,j}$ are independent with probability p each.) The permutation group S_∞ defined earlier acts on the event space $(\Omega, \mathcal{B}_{\max})$ in the obvious manner. Thus for instance $\sigma(\mathcal{B}_I) = \mathcal{B}_{\sigma(I)}$ for all $\sigma \in S_\infty$ and $I \subseteq \mathbf{N}$.

Next, we need a way to embed every finite graph into the universal space.

Definition 3.2 (Graph universal embedding). Let $m \geq 1$, and let $G^{(m)} = (V^{(m)}, E^{(m)})$ be a finite graph. Let $(\Omega^{(m)}, \mathcal{B}_{\max}^{(m)}, \mathbf{P}^{(m)})$ be the probability space corresponding to the sampling of a countable sequence⁸ of i.i.d. random variables $x_1^{(m)}, x_2^{(m)}, \dots \in V^{(m)}$ sampled independently and uniformly at random⁹. To every sequence $(x_1^{(m)}, x_2^{(m)}, \dots) \in \Omega^{(m)}$ we associate an infinite graph $G_\infty^{(m)} = (\mathbf{N}, E_\infty^{(m)}) \in \Omega$ by setting

$$E_\infty^{(m)} := \{\{i, j\} \in \binom{\mathbf{N}}{2} : \{x_i^{(m)}, x_j^{(m)}\} \in E^{(m)}\};$$

one can think of this as a random lifting of the graph $G^{(m)}$ on $V^{(m)}$ up to an infinite graph $G_\infty^{(m)}$ on the natural numbers \mathbf{N} . This mapping from $\Omega^{(m)}$ to Ω is clearly measurable, since the inverse images of the generating events $A_{i,j}$ in $(\Omega, \mathcal{B}_{\max})$ are simply the events that $\{x_i^{(m)}, x_j^{(m)}\}$ lie in $G^{(m)}$, which are certainly measurable in $\mathcal{B}_{\max}^{(m)}$. This allows us to extend the probability measure $\mathbf{P}^{(m)}$ from $(\mathcal{B}_{\max}^{(m)}, \Omega^{(m)})$ to the product space $(\mathcal{B}_{\max} \times \mathcal{B}_{\max}^{(m)}, \Omega \times \Omega^{(m)})$ in a canonical manner, identifying the events $A_{i,j}$ with the events $\{x_i^{(m)}, x_j^{(m)}\} \in E^{(m)}$. We shall abuse notation and refer to the extended measure also as $\mathbf{P}^{(m)}$.

Remarks 3.3. Now that the space $(\Omega^{(m)}, \mathcal{B}_{\max}^{(m)})$ is infinite, not every event involving the $x_i^{(m)}$ is measurable, however any event which involves only finitely many of the $x_i^{(m)}$ is certainly measurable (and in particular has a well-defined probability). One can view $E_\infty^{(m)}$ as the infinite random graph formed by statistically sampling

⁸This sequence contains the “hidden variables” that will play the role of the parameters $x^{(m)}$ and $\lambda^{(m)}$ in the preceding section. Again, the exact construction of this Wiener-type probability space is not important. The most canonical way to proceed is to let $\Omega^{(m)}$ be the countable product $(V^{(m)})^{\mathbf{N}}$ with the product σ -algebra $\mathcal{B}_{\max}^{(m)}$ and the product uniform probability measure $\mathbf{P}^{(m)}$. A more concrete way would be to identify $V^{(m)}$ with $[n^{(m)}] = \{1, \dots, n^{(m)}\}$ by appropriate labeling, set $\Omega^{(m)}$ to be the unit interval $[0, 1] := \{x : 0 \leq x < 1\}$, let $\mathcal{B}_{\max}^{(m)}$ be the Borel σ -algebra, $\mathbf{P}^{(m)}$ be Lebesgue measure, and let x_j be the j^{th} digit in the base- $n^{(m)}$ expansion of x (rounding down when a terminating decimal occurs).

⁹Of course for any fixed m there will be infinitely many repetitions among these $x_i^{(m)}$ since $V^{(m)}$ is finite, but in practice we are interested in taking limits in which $|V^{(m)}| \rightarrow \infty$, and so these collisions will become asymptotically negligible.

of the original finite (and deterministic) graph $E^{(m)}$. This is a convenient way to convert arbitrary graphs, on arbitrary vertex sets, to a fixed universal (random) graph on a fixed universal vertex set, in this case the natural numbers \mathbf{N} . The random graph $E_\infty^{(m)}$ turns out to capture all the relevant features we require of the original graph; for instance, the probability that $E_\infty^{(m)}$ lies in the event $A_{1,2}$ is essentially¹⁰ the edge density of $E^{(m)}$, while the probability that $E_\infty^{(m)}$ lies in $A_{1,2} \wedge A_{2,3} \wedge A_{3,1}$ is essentially the triangle density of $E^{(m)}$, and so forth. On the other hand, it suppresses irrelevant features such as what the labels of the original vertex set $V^{(m)}$ were; in particular, applying a graph isomorphism to $E^{(m)}$ does not affect the probability distribution of $E_\infty^{(m)}$ at all. More generally, we observe the permutation invariance

$$\mathbf{P}^{(m)}(\sigma E) = \mathbf{P}^{(m)}(E) \text{ for all } E \in \mathcal{B}_{\max}, \sigma \in S_\infty \quad (4)$$

which can be verified by first checking on regular events E (i.e. finite boolean combinations of the $A_{i,j}$) and then extending as in the proof of the Furstenberg correspondence principle.

Once again, we can view the random graph $G_\infty^{(m)}$ as being generated¹¹ by “hidden variables” $x_1^{(m)}, x_2^{(m)}, \dots$. As before, we wish to “forget” these hidden variables and pass to a limit. This can be achieved as follows.

Proposition 3.4 (Graph correspondence principle). *For every $m \geq 1$, let $G^{(m)} = (V^{(m)}, E^{(m)})$ be a finite undirected graph, and let $\mathbf{P}^{(m)}$ be as in Definition 3.2. Then there exists a subsequence $0 < m_1 < m_2 < \dots$ of m , and a probability measure $\mathbf{P}^{(\infty)}$ on the graph universal space $(\Omega, \mathcal{B}_{\max})$, such that we have the weak convergence property (1) and the permutation invariance property*

$$\mathbf{P}^{(\infty)}(\sigma E) = \mathbf{P}^{(\infty)}(E) \text{ for all } E \in \mathcal{B}_{\max}, \sigma \in S_\infty. \quad (5)$$

Proof The algebra \mathcal{B}_{reg} is countable. Thus the existence of the weak limit $\mathbf{P}^{(m_i)}$ follows from Lemma A.15. From (4) we can deduce (5) by arguing exactly as in the Furstenberg correspondence principle. ■

So far, the permutation group S_∞ has played the same role for graphs as the integer group \mathbf{Z} played for sets of integers. However, the permutation group is significantly more “mixing”, which allows us to immediately “regularise” the system obtained in Proposition 3.4:

Lemma 3.5 (Infinitary regularity lemma). *Let \mathbf{P} be a probability measure on the graph universal space $(\Omega, \mathcal{B}_{\max})$ which is permutation-invariant in the sense of (5). Then for any $I, I_1, \dots, I_l \subset \mathbf{N}$ with $I \cap I_1 \cap \dots \cap I_l$ infinite, the factors \mathcal{B}_I and $\bigvee_{i=1}^l \mathcal{B}_{I_i}$ are relatively independent conditioning on $\bigvee_{i=1}^l \mathcal{B}_{I \cap I_i}$, with respect to this probability measure \mathbf{P} . (See Appendix A for a definition of relative independence.)*

¹⁰We say “essentially” because there is a slight error term coming from the event that $x_1^{(m)} = x_2^{(m)}$. However this error will become negligible in limits for which $|V^{(m)}| \rightarrow \infty$.

¹¹This is of course the perspective taken in property testing. It is not surprising that the Szemerédi regularity lemma plays a crucial role in that theory also; see [2]. Indeed, this argument suggests that an infinitary approach to property testing theory is possible.

This result is the infinitary analogue of the Szemerédi regularity lemma, and will play a crucial role in establishing the proof of the triangle removal lemma (and later, the hypergraph removal lemma) in subsequent sections.

Proof Fix I, I_1, \dots, I_l . We may assume $l \geq 1$ since the claim is trivial when $l = 0$. To show that \mathcal{B}_I and $\bigvee_{i=1}^l \mathcal{B}_{I_i}$ are relatively independent conditioning on $\bigvee_{i=1}^l \mathcal{B}_{I \cap I_i}$ with respect to the probability measure $\mathbf{P}^{(\infty)}$, it suffices by Lemma A.26 to show that

$$\|\mathbf{P}(E_I | \bigvee_{i=1}^l \mathcal{B}_{I_i})\|_{L^2} = \|\mathbf{P}(E_I | \bigvee_{i=1}^l \mathcal{B}_{I \cap I_i})\|_{L^2}$$

for all $E_I \in \mathcal{B}_I$. By Lemma A.18 and limiting arguments we may assume without loss of generality that E_I is regular. In particular we have $E_I \in \mathcal{B}_{I'}$ for some finite subset I' of I . By Corollary A.20 and a limiting argument we may assume that the set I has an infinite complement. By another such limiting argument we can also assume that $I_i \setminus I$ is finite for all i .

Let A be the infinite set $A := (I \cap I_1 \cap \dots \cap I_l) \setminus I'$, and let B be the finite set $B := \bigcup_{i=1}^l I_i \setminus I$. Then we can find a permutation σ be a permutation which maps A to $A \cup B$ bijectively, but is constant on $I \setminus A$, and in particular fixes I' . Thus σ also fixes E_I , and maps $I \cap I_i$ to $(I \cap I_i) \cup B$. Thus

$$\|\mathbf{P}(E_I | \bigvee_{i=1}^l \mathcal{B}_{(I \cap I_i) \cup B})\|_{L^2} = \|\mathbf{P}(E_I | \bigvee_{i=1}^l \mathcal{B}_{I \cap I_i})\|_{L^2}.$$

But as \mathcal{B}_{I_i} lies between $\mathcal{B}_{I \cap I_i}$ and $\mathcal{B}_{(I \cap I_i) \cup B}$, the claim now follows from Lemma A.12. \blacksquare

Remark 3.6. The above proof of the regularity lemma is short but perhaps a bit opaque. Let us informally discuss a special case of this lemma, namely that the events $A_{1,3}$ and $A_{2,3}$ are relatively independent conditioning on $\mathcal{B}_{\{3,4,5,\dots\}}$; this is a special case of the situation where $I = \mathbf{N} \setminus \{2\}$, $I_1 = \mathbf{N} \setminus \{1\}$, and $l = 1$. Passing back to the finite graph setting (by working with the probability measures $\mathbf{P}^{(m)}$ from Proposition 3.4), this claim may seem puzzling at first, because the events $\{x_1^{(m)}, x_3^{(m)}\} \in E^{(m)}$ and $\{x_2^{(m)}, x_3^{(m)}\} \in E^{(m)}$ can certainly be correlated; indeed, whenever $x_3^{(m)}$ has high degree, then both events occur with high probability, and when it has low degree, both events occur with low probability. However, if one can somehow learn the degree of $x_3^{(m)}$, then these two events become relatively independent *conditioning on the degree of $x_3^{(m)}$* . And now the purpose of the factor $\mathcal{B}_{\{3,4,5,\dots\}}$ becomes clear; by “polling” many additional vertices $x_4^{(m)}, x_5^{(m)}, \dots, x_N^{(m)}$ and measuring the connectivity of $x_3^{(m)}$ with all of these additional vertices, we can obtain a statistical prediction for the degree of $x_3^{(m)}$, whose accuracy and confidence level become almost surely perfect in the asymptotic limit $N \rightarrow \infty$. More generally, it turns out that by polling the interconnectivity of vertices in the infinite set $I \cap I_i$ for $i = 1, \dots, l$ one can obtain an almost surely perfectly accurate prediction of all the “common information” held between an event in \mathcal{B}_I and an event in $\bigvee_{i=1}^l \mathcal{B}_{I_i}$. Let us illustrate this with one further example, namely the relative independence of $A_{1,2}$ and $A_{2,3} \wedge A_{1,3}$ conditioning on $\mathcal{B}_{\{1,4,5,\dots\}} \vee \mathcal{B}_{\{2,4,5,\dots\}}$; this corresponds

to the case $I = \mathbf{N} \setminus \{3\}$, $l = 2$, and $I_i = \mathbf{N} \setminus \{i\}$ for $i = 1, 2$. We are asking for the events $\{x_1^{(m)}, x_2^{(m)}\} \in E^{(m)}$ and $\{x_1^{(m)}, x_3^{(m)}\}, \{x_2^{(m)}, x_3^{(m)}\} \in E^{(m)}$ to become relatively independent once we sample all the connectivity information between $x_2^{(m)}$ and $x_4^{(m)}, x_5^{(m)}, \dots$, and between $x_3^{(m)}$ and $x_4^{(m)}, x_5^{(m)}, \dots$. To see how this will work, observe that while the two events in question will not be unconditionally independent in general, they will become conditionally independent once the number of paths of length two connecting $x_1^{(m)}$ and $x_2^{(m)}$ are known, since upon freezing $x_1^{(m)}$ and $x_2^{(m)}$ this determines the probability that the independent variable $x_3^{(m)}$ will satisfy the latter event $\{x_1^{(m)}, x_3^{(m)}\}, \{x_2^{(m)}, x_3^{(m)}\} \in E^{(m)}$; since $x_3^{(m)}$ does not affect the former event $\{x_1^{(m)}, x_2^{(m)}\} \in E^{(m)}$, we obtain relative independence. But the number of paths of length two can be determined statistically, by counting the proportion of $j \in \{4, 5, \dots\}$ for which $\{x_1^{(m)}, x_j^{(m)}\}$ and $\{x_2^{(m)}, x_j^{(m)}\}$ both lie in $E^{(m)}$. This lies in the factor $\mathcal{B}_{\{1,4,5,\dots\}} \vee \mathcal{B}_{\{2,4,5,\dots\}}$ and is the reason for the conditional independence¹².

Remark 3.7. Similar correspondence principles exist for bipartite graphs, directed graphs, and multicolored graphs (where the color set is kept independent of m), and so forth; for instance, a generalisation to tripartite graphs is sketched out in Appendix B. We will not pursue the other generalisations here as they are rather minor, though we will consider a hypergraph extension of this principle in Section 7.

4. AN INFINITARY PROOF OF THE TRIANGLE REMOVAL LEMMA

Let us now apply the above correspondence principle to obtain the following triangle-removal lemma of Ruzsa and Szemerédi:

Lemma 4.1 (Triangle removal lemma). [24] *Let $G = (V, E)$ be an undirected graph with $|V| = n$ vertices. Suppose that G contains fewer than δn^3 triangles for some $0 < \delta \leq 1$, or more precisely*

$$|\{(x_1, x_2, x_3) \in V^3 : \{x_1, x_2\}, \{x_2, x_3\}, \{x_3, x_1\} \in E\}| \leq \delta n^3.$$

Then it is possible to delete $o_{\delta \rightarrow 0}(n^2)$ edges from G to create a graph G' which is triangle-free. Here $o_{\delta \rightarrow 0}(n^2)$ denotes a quantity, which when divided by n^2 , goes to zero as $\delta \rightarrow 0$, uniformly in n .

Previous to this paper, the only known proof of lemma proceeded via the *Szemerédi regularity lemma* [26]. It can be used among other things to imply the $k = 3$ case of

¹²There is another way of viewing this, namely that each vertex $x_j^{(m)}$ induces a partition of the $x_1^{(m)}$ and $x_2^{(m)}$ vertex sets, by dividing them into those vertices which are connected to $x_j^{(m)}$ in $E^{(m)}$ and those that are not. Letting j vary in $\{4, 5, \dots, N\}$ one obtains a partition of these vertex classes which behaves increasingly like the partitions created by the Szemerédi regularity lemma as $N \rightarrow \infty$, in the sense that the graph between the $x_1^{(m)}$ and $x_2^{(m)}$ becomes increasingly “ ε -regular” relative to this partition; the ε -regularity is closely related to the relative independence properties discussed here. We will however not pursue this approach as it becomes somewhat complicated when we move to the hypergraph setting, whereas the techniques we present here carries over to hypergraphs with virtually no changes.

Szemerédi's theorem (Theorem 2.1). Based on this connection, it is natural to ask whether there is an infinitary analogue of this lemma, similarly to how Theorem 2.2 is an infinitary counterpart to Theorem 2.1. We shall deduce it from the following substantially stronger infinitary statement.

If J is a set, we define an *downset* \mathbf{i} in J to be any collection of subsets e of J with the property that whenever $e \in \mathbf{i}$ and $e' \subseteq e$, then $e' \in \mathbf{i}$ also. In particular, downsets are automatically closed under intersection.

Theorem 4.2 (Hypergraph removal lemma, infinitary version). *Let $(\Omega, \mathcal{B}_{\max}, \mathbf{P})$ be a probability space, and let $\mathcal{B}_{\text{reg}} \subseteq \mathcal{B}_{\max}$ be an algebra. Let J be a finite set, and let \mathbf{i}_{\max} be an downset in J . For each $e \in \mathbf{i}_{\max}$ let \mathcal{B}_e be a factor of \mathcal{B}_{\max} with the following properties:*

- (Regularisability) *Each of the factors \mathcal{B}_e is generated by countably many events from \mathcal{B}_{reg} .*
- (Nesting) *If $e, e' \in \mathbf{i}_{\max}$ are such that $e \subseteq e'$, then \mathcal{B}_e is a factor of $\mathcal{B}_{e'}$.*
- (Independence) *If $e, e_1, \dots, e_l \in \mathbf{i}_{\max}$, then the factors \mathcal{B}_e and $\bigvee_{i=1}^l \mathcal{B}_{e_i}$ are relatively independent conditioning on $\bigvee_{i=1}^l \mathcal{B}_{e \cap e_i}$.*

For each $e \in \mathbf{i}_{\max}$, let E_e be an event in \mathcal{B}_e such that

$$\mathbf{P}\left(\bigwedge_{e \in \mathbf{i}_{\max}} E_e\right) = 0.$$

Then for any $\varepsilon > 0$, there exist events $F_e \in \mathcal{B}_e \cap \mathcal{B}_{\text{reg}}$ for all $e \in \mathbf{i}_{\max}$ such that

$$\mathbf{P}(E_e \setminus F_e) \leq \varepsilon \text{ for all } e \in \mathbf{i}_{\max}$$

and

$$\bigwedge_{e \in \mathbf{i}_{\max}} F_e = \emptyset.$$

We will prove this rather strange-looking proposition in Section 6. For the purposes of proving the triangle removal lemma, we will only need this lemma in the special case when $J = \{1, 2, 3\}$, when $\mathbf{i}_{\max} := \{e : |e| \leq 2\}$, and when $A_e = \Omega$ for all $e \neq \{1, 2\}, \{2, 3\}, \{3, 1\}$. However, the lemma is not that much more difficult to prove in the general case¹³, and it will rather easily yield a hypergraph generalisation of the triangle removal lemma, so we retain the proposition in the general form. The hypothesis $\mathbf{P}(\bigwedge_{e \in \mathbf{i}_{\max}} E_e) = 0$ is the analogue in Lemma 4.1 of the hypothesis that G has few triangles, while the conclusion $\bigwedge_{e \in \mathbf{i}_{\max}} F_e = \emptyset$ is the analogue of the conclusion that the modified graph G' is triangle-free.

¹³This is in stark contrast to the finitary situation, in which the hypergraph removal lemma is significantly more difficult than the triangle removal lemma. This is ultimately because of the need in the finitary hypergraph setting to constantly play off epsilons of different sizes against each other; see [19], [20], [22], [23], [12], [30] for some examples of this. However in the infinitary asymptotic limit, most of the epsilons have disappeared or at least been confined to individual lemmas where they do not interact with other epsilons. This simplifies the proof significantly, albeit at the cost of working in an infinitary setting as opposed to a finitary one. In the converse direction, note the proliferation of epsilons in [29] when Furstenberg's proof of Szemerédi's theorem is transferred from the infinitary setting to the finitary one.

Proof [of Lemma 4.1 assuming Theorem 4.2] Suppose for contradiction that Lemma 4.1 failed. Then we can find an $0 < \eta \leq 1$ and sequence $n^{(m)}$ of integers, a sequence of graphs $G^{(m)} = (V^{(m)}, E^{(m)})$ with $|V^{(m)}| = n^{(m)}$, such that the $G^{(m)}$ have asymptotically vanishing number of triangles,

$$\lim_{m \rightarrow \infty} \frac{1}{(n^{(m)})^3} |\{(x_1, x_2, x_3) \in (V^{(m)})^3 : \{x_1, x_2\}, \{x_2, x_3\}, \{x_3, x_1\} \in E^{(m)}\}| = 0 \quad (6)$$

but such that each of the $G^{(m)}$ cannot be made triangle-free without deleting at least $\eta(n^{(m)})^2$ edges. (One could make the decay rate in (6) more rapid, as in the proof of Theorem 2.1, but we will find no need to do so here.) In particular, $G^{(m)}$ contains at least one triangle, and hence the expression inside the limit of (6) is at least $1/(n^{(m)})^3$. This implies that

$$n^{(m)} \rightarrow \infty \text{ as } m \rightarrow \infty. \quad (7)$$

Now let $(\Omega, \mathcal{B}_{\max})$ be the graph universal space introduced in Definition 3.1, with the attendant events $A_{n,m}$, regular algebra \mathcal{B}_{reg} , factors \mathcal{B}_I , and S_∞ group action. Let $\mathbf{P}^{(m)}$ be the probability measure on $(\Omega \times \Omega^{(m)}, \mathcal{B}_{\max} \times \mathcal{B}_{\max}^{(m)})$ defined in Definition 3.2, and let $\mathbf{P}^{(\infty)}$ be a limiting measure as constructed in the graph correspondence principle (Proposition 3.4). From (6) we have

$$\lim_{m \rightarrow \infty} \mathbf{P}^{(m)}(A_{1,2} \wedge A_{2,3} \wedge A_{3,1}) = 0$$

and hence by (1)

$$\mathbf{P}^{(\infty)}(A_{1,2} \wedge A_{2,3} \wedge A_{3,1}) = 0.$$

We will apply Theorem 4.2 on the universal space $(\Omega, \mathcal{B}_{\max}, \mathbf{P}^{(\infty)})$ with $J := \{1, 2, 3\}$, $\mathbf{i}_{\max} := \{e : |e| \leq 2\}$, E_e set equal to $A_{i,j}$ if $e = \{i, j\}$ for some $ij = 12, 23, 31$, and $E_e = \Omega$ otherwise, and with \mathcal{B}_e set equal to $\mathcal{B}_{e \cup \{4, 5, \dots\}}$ for all $e \in \mathbf{i}_{\max}$. The nesting and regularisability properties required for Theorem 4.2 are obvious, while the independence properties follow from Lemma 3.5. We can thus invoke the theorem and find regular events $F_e \in \mathcal{B}_{e \cup \{4, 5, \dots\}} \cap \mathcal{B}_{\text{reg}}$ for $e \in \mathbf{i}_{\max}$ with

$$\mathbf{P}^{(\infty)}(E_e \setminus F_e) < \eta/100 \text{ for } e \in \mathbf{i}_{\max} \quad (8)$$

such that

$$\bigwedge_{e \in \mathbf{i}_{\max}} F_e = \emptyset. \quad (9)$$

It is convenient to eliminate the lower order components $\emptyset, \{1\}, \{2\}, \{3\}$ of the ideal \mathbf{i}_{\max} . For $ij = 12, 23, 31$, define $F'_{i,j} := F_{\{i,j\}} \wedge F_{\{i\}} \wedge F_{\{j\}} \wedge F_\emptyset$. Then the $F'_{i,j}$ are regular, and (by monotonicity) we have $F'_{i,j} \in \mathcal{B}_{\{i,j,4,5,\dots\}}$. From (8) and the choice of the E_e we have

$$\mathbf{P}^{(\infty)}(A_{i,j} \setminus F'_{i,j}) < \eta/10 \text{ for } ij = 12, 23, 31 \quad (10)$$

while from (9) we have

$$F'_{1,2} \wedge F'_{2,3} \wedge F'_{3,1} = \emptyset. \quad (11)$$

Now we reinstate the “hidden variables” $x_1^{(m)}, x_2^{(m)}, \dots$ by viewing $\mathbf{P}^{(m)}$ once again as a probability measure on the product space $(\Omega \times \Omega^{(m)}, \mathcal{B}_{\max} \times \mathcal{B}_{\max}^{(m)})$; in particular $A_{i,j}$ is now identified with the event that $\{x_i^{(m)}, x_j^{(m)}\}$ lies in the graph $G^{(m)}$. Now because $A_{i,j}$ and $F'_{i,j}$ are regular, the quantity $\mathbf{P}^{(m)}(A_{i,j} \setminus F'_{i,j})$ is the probability of an event involving only finitely many of the random vertices $x_i^{(m)}$ of $V^{(m)}$; let us say that it involves only the vertices $x_1^{(m)}, \dots, x_N^{(m)}$ (note that N will be independent of m , depending only on the complexity of the event $F'_{i,j}$). By increasing N if necessary we may assume $N > 3$. Recall that $F'_{i,j}$ depends only on the vertices $x_i^{(m)}, x_j^{(m)}$ and $x_4^{(m)}, \dots, x_N^{(m)}$. For any fixed values of $x_4^{(m)}, \dots, x_N^{(m)}$, let us say that a vertex pair $\{x, y\} \subset V^{(m)}$ is *good* if for each $ij = 12, 23, 31$, the event $F'_{i,j}$ holds true whenever x, y are substituted for either $x_i^{(m)}, x_j^{(m)}$ or $x_j^{(m)}, x_i^{(m)}$. Now define the random subgraph $(G')^{(m)} = (V^{(m)}, (E')^{(m)})$ of $G^{(m)} = (V^{(m)}, E^{(m)})$ by setting $(E')^{(m)}$ to be all the good pairs $\{x, y\}$ in $E^{(m)}$; this graph depends on the random variables $x_4^{(m)}, \dots, x_N^{(m)}$. From (10) we see that

$$\mathbf{E}|E^{(m)} \setminus (E')^{(m)}| < \eta(N^{(m)})^2.$$

Also, we observe that regardless of the values of $x_3^{(m)}, \dots, x_N^{(m)}$, the graph $(G')^{(m)}$ almost surely cannot contain any triangles, as this would contradict (9). But by the pigeonhole principle we can find a deterministic representative of the random graph $(G')^{(m)}$ for which

$$|E^{(m)} \setminus (E')^{(m)}| < \eta(N^{(m)})^2,$$

and so we have made $G^{(m)}$ triangle-free by removing fewer than $\eta(N^{(m)})^2$ edges, a contradiction that establishes Lemma 4.1. \blacksquare

Remark 4.3. The same arguments in fact give a *subgraph removal lemma*, in which the triangle is replaced by another fixed subgraph. The proof is the same, it is only the downset \mathbf{i}_{\max} (and some minor numerical factors in the argument) which change significantly. But in all these cases, the elements in the downset will only have cardinality at most two. We will not give the details here since they will be subsumed by the hypergraph removal lemma in Theorem 8.1. The higher order cases of Theorem 4.2, involving sets e of three or more elements, do not actually get used in graph theory (which is ultimately only concerned with finite boolean combinations of relations that involve at most two vertices at a time), and are only of importance for hypergraph theory (in which one must now consider combinations of relations, each of which involve three or more vertices).

As observed in [30], there is a slightly stronger version of the triangle removal lemma which gives some further complexity information on G' , at the expense of conceding that G' need not be a subgraph of G . More precisely, we have

Lemma 4.4 (Strong triangle removal lemma). *Let $G = (V, E)$ be an undirected graph with $|V| = n$ vertices. Suppose that G contains fewer than δn^3 triangles for some $0 < \delta \leq 1$. Then one can find a triangle-free graph G' with $G \setminus G'$ containing fewer than $o_{\delta \rightarrow 0}(n^2)$ edges. Furthermore, there exists a partition of V into $O_{\delta}(1)$ components, such that when restricted to the edges joining any two of these*

partitions (which could be equal), then G' is either a complete graph or an empty graph.

This stronger version of the lemma is a by-product of the usual proof of Lemma 4.1, as the graph G' is constructed by excluding certain bad pairs of Szemerédi cells from the graph G . It turns out that the infinitary approach can also yield this stronger lemma without much difficulty.

Proof We again argue by contradiction. But this time, the contradiction hypothesis yields a more complicated statement. More precisely, if Lemma 4.4 failed, then we can find an $0 < \eta \leq 1$ and sequence $n^{(m)}$ of integers, a sequence of graphs $G^{(m)} = (V^{(m)}, E^{(m)})$ with $|V^{(m)}| = n^{(m)}$, and a sequence $M^{(m)}$ tending to infinity as $m \rightarrow \infty$, such that (6) holds, but such that there does not exist any triangle-free graph G' for which $G \setminus G'$ has fewer than $\eta(n^{(m)})^2$ edges, and for which there exists a partition of $V^{(m)}$ into $M^{(m)}$ or fewer components, such that when restricted to the edges joining any two of these cells, G' is either the complete graph or the empty graph.

We now repeat all the arguments used to prove Lemma 4.1, until we get to the point where we have created regular events $F'_{i,j} \in \mathcal{B}_{\{i,j,4,5,\dots\}}$ obeying (10) and (11). Now we insert an additional step to lower the complexity of the events $F'_{i,j}$. Observe that $\mathcal{B}_{\{i,j,4,5,\dots\}}$ is generated by the factor $\mathcal{B}_{\{i,4,5,\dots\}} \vee \mathcal{B}_{\{j,4,5,\dots\}}$, together with the additional event $A_{i,j}$. Thus we can write $F'_{i,j} \wedge A_{i,j} = F''_{i,j} \wedge A_{i,j}$ for some $F''_{i,j} \in \mathcal{B}_{\{i,4,5,\dots\}} \vee \mathcal{B}_{\{j,4,5,\dots\}}$. From (10) we have

$$\mathbf{P}^{(\infty)}(A_{i,j} \setminus F''_{i,j}) < \eta/10 \text{ for } ij = 12, 23, 31.$$

Now we argue that we still have the analogue of (11), namely

$$F''_{1,2} \wedge F''_{2,3} \wedge F''_{3,1} = \emptyset. \quad (12)$$

From (11) we already have

$$(F''_{1,2} \wedge F''_{2,3} \wedge F''_{3,1}) \cap (A_{1,2} \wedge A_{2,3} \wedge A_{3,1}) = \emptyset. \quad (13)$$

But the regular event $F''_{1,2} \wedge F''_{2,3} \wedge F''_{3,1}$ is a boolean combination of finitely many events $A_{i,j}$, where at most one of the i, j lie in $\{1, 2, 3\}$. In other words, this combination does not involve $A_{1,2} \wedge A_{2,3} \wedge A_{3,1}$. Thus if (12) failed, so that there was an infinite graph (\mathbf{N}, E_∞) lying in $F''_{1,2} \wedge F''_{2,3} \wedge F''_{3,1}$, we could modify the graph E_∞ on the edges $\{1, 2\}, \{2, 3\}, \{3, 1\}$ so that it also lies in the set in (13), a contradiction.

To summarise, we can safely replace $F'_{i,j}$ by the lower complexity event $F''_{i,j}$. Now we continue the argument in the proof of Lemma 4.1 with this replacement, but define the edges of $(G')^{(m)}$ to be all the good pairs $\{x, y\}$ in $V^{(m)}$ rather than in $E^{(m)}$. This means that $(G')^{(m)}$ is no longer a subgraph of $(G)^{(m)}$, but the property of being good is determined entirely by the regular events $F''_{i,j}$, which in turn only involve finitely many events $A_{i,j}$ with at most one of the i, j lying in $\{1, 2, 3\}$. Inspecting the definition of a good pair, we see that for fixed $x_4^{(m)}, \dots, x_N^{(m)}$ for N sufficiently large, the graph $(G')^{(m)}$ has bounded complexity, in the sense that there is a partition of $V^{(m)}$ into M cells, for some M depending only on the $F''_{i,j}$, such that

when restricted to the edges joining any two of these cells, G' is either the complete graph or the empty graph. But for m sufficiently large we have $M^{(m)} > M$, and so we attain the same contradiction as before. \blacksquare

5. THE UNIFORM INTERSECTION PROPERTY

We now build the machinery necessary to prove the infinitary hypergraph removal lemma (Theorem 4.2). Again, we will be motivated by the example from ergodic theory. In Furstenberg's proof [7], [10], [8] of the Furstenberg recurrence theorem (Theorem 2.2), the proof proceeded by a kind of induction on factors. Let us say that a factor \mathcal{B} of \mathcal{B}_{\max} obeys the *uniform multiple recurrence (UMR) property* if the conclusion of Theorem 2.2 holds true whenever $E \in \mathcal{B}$ and $\mathbf{P}(E) > 0$. Thus for instance the trivial factor $\{\emptyset, \Omega\}$ has the UMR property. One then shows that the UMR property is preserved under three operations: weakly mixing extensions; limits of chains; and compact (or finite rank) extensions. An application of Zorn's lemma¹⁴ then allows one to conclude that the maximal factor \mathcal{B}_{\max} obeys the UMR property.

We will adopt a similar strategy here, based around a certain property of families of factors which we call the *uniform intersection property (UIP)*. This property is again trivial for very small families, and will be preserved under the same three operations of weakly mixing extensions, limits of chains, and finite rank extensions. Because of the finiteness of J in Theorem 4.2, we will only need to apply these operations finitely often, and will not require Zorn's lemma. However it does seem likely that there are extensions of this theorem to the case when J is infinite, and (more interestingly) to the case where the sets e in \mathbf{i}_{\max} can be unbounded or even countably infinite. We will not pursue this matter here.

We begin by stating the UIP.

Definition 5.1 (Uniform intersection property). Let $(\Omega, \mathcal{B}_{\max}, \mathbf{P})$ be a probability space, and let \mathcal{B}_{reg} be an algebra in \mathcal{B}_{\max} . We say that a tuple $(\mathcal{B}_i)_{i \in I}$ of factors has the *uniform intersection property (UIP)* if the following holds: given any tuple $(E_i)_{i \in I}$ of events $E_i \in \mathcal{B}_i$ with $\mathbf{P}(\bigwedge_{i \in I} E_i) = 0$, and given any $\varepsilon > 0$, there exists a tuple $(F_i)_{i \in I}$ of regular events $F_i \in \mathcal{B}_i \wedge \mathcal{B}_{\text{reg}}$ with $\mathbf{P}(E_i \setminus F_i) \leq \varepsilon$ for each $i \in I$ such that $\bigwedge_{i \in I} F_i = \emptyset$.

Remark 5.2. Roughly speaking, the UIP asserts that if events E_i from separate factors \mathcal{B}_i have a null intersection, then this fact can be almost entirely “explained” by *regular* events $F_i \in \mathcal{B}_i$ which have *empty* intersection. Thus, for instance, the conclusion of Theorem 4.2 is simply that the tuple $(\mathcal{B}_I)_{I \in \mathbf{i}_{\max}}$ obeys the UIP.

¹⁴Actually, to establish Theorem 2.2 for a fixed k , one only needs to apply the limits-of-chains step a finite number of times depending on k , at which point one reaches a factor which is *characteristic* for the maximal factor \mathcal{B}_{\max} , at which point one can jump directly to \mathcal{B}_{\max} without using Zorn's lemma. Thus the proof of the Furstenberg recurrence theorem does not actually require the axiom of choice, and indeed the original proof in [7] did not use this axiom.

Before continuing, let us illustrate the UIP with a few simple examples. All of these examples take place in some probability space $(\Omega, \mathcal{B}_{\max}, \mathbf{P})$ with an algebra \mathcal{B}_{reg} of regular events. We say that a factor \mathcal{B} is *regularisable* if it can be generated by at most countably many regular events.

Example 5.3. The empty tuple $()$ obeys the UIP in a vacuous sense (the hypothesis $\mathbf{P}(\bigwedge_{i \in I} E_i) = 0$ is impossible to satisfy).

Example 5.4. Let \mathcal{B} be a factor. Then the singleton tuple (\mathcal{B}) trivially has the UIP (indeed one can even take $\varepsilon = 0$ and $F_{\mathcal{B}} = \emptyset$ in this case).

Example 5.5. Let \mathcal{B} be a regularisable factor. Then the 2-tuple $(\mathcal{B}, \mathcal{B})$ has the UIP. Indeed, if $E, E' \in \mathcal{B}$ were such that $\mathbf{P}(E \wedge E') = 0$, then from Lemma A.19 one can find regular events $\tilde{F}, \tilde{F}' \in \mathcal{B}$ which are $\varepsilon/3$ -close to E, E' respectively. By the triangle inequality this implies that $\mathbf{P}(\tilde{F} \wedge \tilde{F}') \leq 2\varepsilon/3$. If we set $F := \tilde{F} \setminus \tilde{F}'$ and $F' := \tilde{F}' \setminus \tilde{F}$ then we see that F, F' are regular events in \mathcal{B} with $F \wedge F' = \emptyset$, while from the triangle inequality $\mathbf{P}(E \setminus F), \mathbf{P}(E' \setminus F') \leq \varepsilon$, and the claim follows. For a generalization of this argument, see Lemma 5.11 below.

Example 5.6. The trivial factor $\{\emptyset, \Omega\}$ has no impact on the UIP. More precisely, a tuple $(\mathcal{B}_i)_{i \in I}$ obeys the UIP if and only if $(\mathcal{B}_i)_{i \in I} \uplus (\{\emptyset, \Omega\})$ also obeys the UIP, where we use \uplus to denote the concatenation of tuples.

Example 5.7. Let $(\mathcal{B}_1, \dots, \mathcal{B}_l)$ be a tuple of jointly independent factors. Then $(\mathcal{B}_1, \dots, \mathcal{B}_l)$ has the UIP. Indeed, if $E_j \in \mathcal{B}_j$ then by joint independence we have $\mathbf{P}(\bigwedge_{1 \leq j \leq l} E_j) = \prod_{j=1}^l \mathbf{P}(E_j)$. Thus if $\bigwedge_{1 \leq j \leq l} E_j$ is a null event, then one of the E_j , say E_{j_0} , must also be a null event. The claim then follows by letting $F_{j_0} = \emptyset$ and letting all the other F_j be the full event Ω . For a more sophisticated version of this argument, see Lemma 5.12 below.

Example 5.8. Let Ω be the unit interval $[0, 1]$ with Lebesgue measure, let \mathcal{B}_{reg} consist of all the finite unions of intervals (open, closed, or half-open), let \mathcal{B}_1 be the factor generated by the event $E_1 := [0, 1/2]$, and let \mathcal{B}_2 be the factor generated by the event $E_2 := [1/2, 1]$. Then $(\mathcal{B}_1, \mathcal{B}_2)$ does not have the UIP. However if one modifies \mathcal{B}_2 to be the factor generated by $(1/2, 1]$ instead of $[1/2, 1]$, then the UIP is restored. Thus the UIP is sensitive to modification of the underlying factors by null events. (On the other hand, the events E_i themselves can be modified by null events within \mathcal{B}_i without any impact to the UIP.)

Example 5.9. Let $\mathcal{B}_1, \mathcal{B}_2$ be finite σ -algebras, and let \mathcal{B} be another σ -algebra, such that $(\mathcal{B}_1, \mathcal{B}_2, \mathcal{B})$ has the UIP. Then $(\mathcal{B}_1 \vee \mathcal{B}_2, \mathcal{B})$ also has the UIP. To see this, let $E_{12} \in \mathcal{B}_1 \vee \mathcal{B}_2$ and $E \in \mathcal{B}$ be such that $\mathbf{P}(E_{12} \vee E) = 0$. Since \mathcal{B}_1 and \mathcal{B}_2 are finite, we can write E_{12} as the union of M events of the form $E_{1,m} \wedge E_{2,m}$ for $1 \leq m \leq M$ for some finite M and some events $E_{1,m} \in \mathcal{B}_1, E_{2,m} \in \mathcal{B}_2$. From the UIP hypothesis we can find regular events $F_{1,m} \in \mathcal{B}_1, F_{2,m} \in \mathcal{B}_2, F_m \in \mathcal{B}$ with $F_{1,m} \wedge F_{2,m} \wedge F_m = \emptyset$ and $\mathbf{P}(E_{1,m} \setminus F_{1,m}), \mathbf{P}(E_{2,m} \setminus F_{2,m}), \mathbf{P}(E \setminus F_m) \leq \varepsilon/M$. If we then set $F_{12} := \bigvee_{m=1}^M (F_{1,m} \wedge F_{2,m})$ and $F := \bigwedge_{m=1}^M F_m$ then the claim follows. For a generalization of this argument, see Lemma 5.13 below.

Remark 5.10. If $(\mathcal{B}_i)_{i \in I}$ has the UIP, then given any tuple $(E_i)_{i \in I}$ of events $E_i \in \mathcal{B}_i$ such that $\bigwedge_{i \in I} E_i$ is a null event, then there exists a tuple $(G_i)_{i \in I}$ of null events $G_i \in \mathcal{B}_i$ which cover the null event $\bigwedge_{i \in I} E_i$. This follows from applying the UIP

with $\varepsilon = 2^{-n}$ (say) to obtain events $(F_{i,n})_{i \in I}$ with $\mathbf{P}(E_i \setminus F_{i,n}) \leq 2^{-n}$ for all $i \in I$ and $\bigwedge_{i \in I} F_{i,n} = \emptyset$, and then letting G_i be the event E_i holds, but that $F_{i,n}$ fails for infinitely many $n \geq 1$; the claim $\bigwedge_{i \in I} E_i \subseteq \bigvee_{i \in I} G_i$ then follows from the pigeonhole principle, while the claim that G_i is null follows from the Borel-Cantelli lemma. Unfortunately the G_i will in general not be regular, and so this consequence of the UIP, while simple to state, is not useful for applications.

We now develop the general tools that we shall use to deduce the UIP for complex tuples from the UIP for simpler tuples. We first show that repetitions do not affect the UIP so long as the σ -algebra being repeated is regularisable. We use $(\mathcal{B}_i)_{i \in I} \uplus (\mathcal{B}_j)_{j \in J} = (\mathcal{B}_i)_{i \in I \uplus J}$ to denote the concatenation of two tuples, where $I \uplus J$ is the disjoint union of I and J (thus one may have to relabel the index set of I or J in order to define this concatenation).

Lemma 5.11. *Let $(\mathcal{B}_i)_{i \in I}$ be a tuple of σ -algebras, and let \mathcal{B} be a regularisable σ -algebra. Then $(\mathcal{B}_i)_{i \in I} \uplus (\mathcal{B})$ has the UIP if and only if $(\mathcal{B}_i)_{i \in I} \uplus (\mathcal{B}, \mathcal{B})$.*

Proof First suppose that $(\mathcal{B}_i)_{i \in I} \uplus (\mathcal{B}, \mathcal{B})$ has the UIP. Then if $E_i \in \mathcal{B}_i$ and $E \in \mathcal{B}$ are such that $\mathbf{P}(E \wedge \bigwedge_{i \in I} E_i) = 0$, then by the UIP hypothesis (inserting a dummy event Ω for the extra copy of \mathcal{B}) we can find regular events $F_i \in \mathcal{B}_i$ and $F', F'' \in \mathcal{B}$ such that $F' \wedge F'' \wedge \bigwedge_{i \in I} F_i = \emptyset$ and

$$\mathbf{P}(E_i \setminus F_i), \mathbf{P}(E \setminus F'), \mathbf{P}(\overline{F''}) \leq \varepsilon/2 \text{ for all } i \in I.$$

The claim then follows by setting $F := F' \wedge F''$.

Now suppose conversely that $(\mathcal{B}_i)_{i \in I} \uplus (\mathcal{B})$ has the UIP. Then if $E_i \in \mathcal{B}_i$ and $E, E' \in \mathcal{B}$ are such that $\mathbf{P}(E \wedge E' \wedge \bigwedge_{i \in I} E_i) = 0$ then by the UIP hypothesis (replacing E and E' by the single event $E \wedge E'$) one can find regular events $F_i \in \mathcal{B}_i$ and $\tilde{F} \in \mathcal{B}$ such that $\tilde{F} \wedge \bigwedge_{i \in I} F_i = \emptyset$ and

$$\mathbf{P}(E_i \setminus F_i), \mathbf{P}((E \wedge E') \setminus \tilde{F}) \leq \varepsilon/3 \text{ for all } i \in I.$$

Now since \mathcal{B} is regularisable, we see from Lemma A.19 that every event in \mathcal{B} is ε -close to a regular event in \mathcal{B} for any $\varepsilon > 0$. In particular, we can find regular events $\tilde{E}, \tilde{E}' \in \mathcal{B}$ which are $\varepsilon/3$ -close to E and E' respectively. If one then sets $F := (\tilde{E} \setminus \tilde{E}') \vee \tilde{F}$ and $F' := (\tilde{E}' \setminus \tilde{E}) \vee \tilde{F}$, we see from the triangle inequality that $\mathbf{P}(E \setminus F), \mathbf{P}(E' \setminus F') \leq \varepsilon$, and that $F \wedge F' \wedge \bigwedge_{i \in I} F_i = \emptyset$, and the claim follows. ■

Now we give the three major extendability properties of the UIP, under weakly mixing extensions, finite rank extensions, and limits of chains. We begin with the analogue of the weakly mixing extension property, which says that one can extend any member of a tuple without destroying the UIP, as long as the extension is relatively independent of all the other factors in the tuple.

Lemma 5.12 (Weakly mixing extensions). *Let $(\mathcal{B}_i)_{i \in I}$ be a tuple of σ -algebras, and let \mathcal{B} be an additional σ -algebra such that $(\mathcal{B}_i)_{i \in I} \uplus (\mathcal{B})$ obeys the UIP. Let \mathcal{B}' be an extension of \mathcal{B} which is relatively independent of $\bigvee_{i \in I} \mathcal{B}_i$ over \mathcal{B} . Then $(\mathcal{B}_i)_{i \in I} \uplus (\mathcal{B}')$ also obeys the UIP.*

Proof Let $E_i \in \mathcal{B}_i$ for $i \in I$ and $E' \in \mathcal{B}'$ be events such that $\mathbf{P}(E' \wedge \bigwedge_{i \in I} E_i) = 0$. We rewrite this as

$$\mathbf{E}(\mathbf{I}(E') \prod_{i \in I} \mathbf{I}(E_i)) = 0.$$

The first factor is measurable in \mathcal{B}' , while the second factor is measurable in $\bigvee_{i \in I} \mathcal{B}_i$. Since these two σ -algebras are relatively independent over \mathcal{B} , we may use (16) and conclude that $\mathbf{P}(E'|\mathcal{B}) \prod_{i \in I} \mathbf{I}(E_i) = 0$ almost surely. Let $E \in \mathcal{B}$ be the support of $\mathbf{P}(E'|\mathcal{B})$ (which is determined only up to a null event in \mathcal{B}), then $E \wedge \bigwedge_{i \in I} E_i$ is a null event. Applying the UIP hypothesis, we can find regular events $F_i \in \mathcal{B}_i$ for $i \in I$ and a regular event $F \in \mathcal{B}$ such that $\mathbf{P}(E_i \setminus F_i), \mathbf{P}(E \setminus F) \leq \varepsilon$ and

$$F \vee \bigvee_{i \in I} F_i = \emptyset.$$

We then set $F' := F$. We will be done as soon as we check that $\mathbf{P}(E' \setminus F') \leq \varepsilon$, which will follow if we can show that $\mathbf{P}(E' \setminus E) = 0$. But

$$\mathbf{P}(E' \setminus E) = \mathbf{P}(\mathbf{I}(E')(1 - \mathbf{I}(E))) = \mathbf{P}(\mathbf{P}(E'|\mathcal{B})(1 - \mathbf{I}(E)))$$

since $1 - \mathbf{I}(E)$ is \mathcal{B} -measurable. But this vanishes by the definition of E' . \blacksquare

Next we turn to the preservation of the UIP under compact extensions (or more accurately “finite rank extensions”), which assert that one can extend any given element of a tuple by finite factors of other elements in the tuple (destroying those elements in the process).

Lemma 5.13 (Finite rank extensions). *Let \mathcal{B}_0 be a σ -algebra, let $\mathcal{B}_1, \dots, \mathcal{B}_l$ be factors of \mathcal{B}_0 , and let $\mathcal{B}'_1, \dots, \mathcal{B}'_l$ be finite σ -algebras for some $l \geq 1$. Let $(\tilde{\mathcal{B}}_i)_{i \in I}$ be an additional tuple of σ -algebras. Then if $(\tilde{\mathcal{B}}_i)_{i \in I} \uplus (\mathcal{B}_1 \vee \mathcal{B}'_1, \dots, \mathcal{B}_l \vee \mathcal{B}'_l, \mathcal{B}_0)$ has the UIP, then $(\tilde{\mathcal{B}}_i)_{i \in I} \uplus (\mathcal{B}_0 \vee \mathcal{B}'_1 \vee \dots \vee \mathcal{B}'_l)$ also has the UIP.*

Proof Write $\mathcal{B}_* := \mathcal{B}_0 \vee \mathcal{B}'_1 \vee \dots \vee \mathcal{B}'_l$. Let E_* be an event in \mathcal{B}_* , and let $\tilde{E}_i \in \tilde{\mathcal{B}}_i$ for $i \in I$, be such that $\mathbf{P}(E_* \wedge \bigwedge_{i \in I} \tilde{E}_i) = 0$. Since $\mathcal{B}'_1, \dots, \mathcal{B}'_l$ are finite, we can write E_* as the finite union of events

$$E_* = E_{*,1} \vee \dots \vee E_{*,M}$$

for some $M \geq 1$, where each $E_{*,m}$ has the form

$$E_{*,m} = E_{0,m} \wedge E_{1,m} \wedge \dots \wedge E_{l,m}$$

for some events $E_{0,m} \in \mathcal{B}_0$ and $E_{j,m} \in \mathcal{B}'_j$ for $1 \leq j \leq l$. For each $1 \leq m \leq M$, we have $\mathbf{P}(E_{*,m} \wedge \bigwedge_{i \in I} \tilde{E}_i) = 0$ and hence

$$\mathbf{P}(E_{0,m} \wedge E_{1,m} \wedge \dots \wedge E_{l,m} \wedge \bigwedge_{i \in I} \tilde{E}_i) = 0.$$

Observe that $E_{j,m} \in \mathcal{B}_j \vee \mathcal{B}'_j$ for $1 \leq j \leq l$, and hence by the UIP hypothesis we may find regular events $F_{0,m} \in \mathcal{B}_0$, $F_{j,m} \in \mathcal{B}_j \vee \mathcal{B}'_j \subseteq \mathcal{B}_*$ for $1 \leq j \leq l$, and $\tilde{F}_{i,m} \in \tilde{\mathcal{B}}_i$ for $i \in I$ such that

$$\mathbf{P}(E_{0,m} \setminus F_{0,m}), \mathbf{P}(E_{j,m} \setminus F_{j,m}), \mathbf{P}(\tilde{E}_i \setminus \tilde{F}_{i,m}) \leq \frac{\varepsilon}{M(l+1)}$$

for $1 \leq j \leq l$ and $i \in I$, and

$$\bigwedge_{j=1}^l F_{j,m} \wedge F_{0,m} \wedge \bigwedge_{i \in I} \tilde{F}_{i,m} = \emptyset.$$

Thus if we set $F_* := \bigvee_{m=1}^M (\bigwedge_{j=1}^l F_{j,m} \wedge F_{0,m})$ and $\tilde{F}_i := \bigwedge_{m=1}^M \tilde{F}_{i,m}$ for $i \in I$ then $F_* \in \mathcal{B}_*$ and $\tilde{F}_i \in \tilde{\mathcal{B}}_i$ are regular events, and $F_* \wedge \bigwedge_{i \in I} \tilde{F}_i = \emptyset$, and

$$\mathbf{P}(E_* \setminus F_*) \leq \sum_{m=1}^M \mathbf{P}(E_{*,m} \setminus (\bigwedge_{j=1}^l F_{j,m} \wedge F_{0,m})) \leq M(l+1) \frac{\varepsilon}{M(l+1)} = \varepsilon$$

and

$$\mathbf{P}(\tilde{E}_i \setminus \tilde{F}_i) \leq \sum_{m=1}^M \mathbf{P}(\tilde{E}_i \setminus \tilde{F}_{i,m}) \leq M \frac{\varepsilon}{M(l+1)} \leq \varepsilon$$

and the claim follows. \blacksquare

Finally, we consider the preservation of the UIP under limits of chains assuming a certain relative independence property.

Lemma 5.14 (Limits of chains). *Let A be a totally ordered set, let I be a finite index set, and for each $\alpha \in A$ let $(\mathcal{B}_{\alpha,i})_{i \in I}$ be a tuple of σ -algebras obeying the UIP, which is increasing in the sense that $\mathcal{B}_{\alpha,i}$ is a factor of $\mathcal{B}_{\beta,i}$ whenever $\alpha < \beta$ and $i \in I$. Let $\mathcal{B}_i := \bigvee_{\alpha \in A} \mathcal{B}_{\alpha,i}$, and suppose that whenever $i \in I$ and $\alpha \in A$, the σ -algebras \mathcal{B}_i and $\bigvee_{j \in I \setminus \{i\}} \mathcal{B}_{\alpha,j}$ are relatively independent over $\mathcal{B}_{\alpha,i}$. Then the tuple $(\mathcal{B}_i)_{i \in I}$ also obeys the UIP.*

Proof Let $E_i \in \mathcal{B}_i$ for $i \in I$ be such that $\mathbf{P}(\bigwedge_{i \in I} E_i) = 0$. From Corollary A.20, we see that for each $i \in I$ there exists an $\alpha \in A$ such that E_i is $\varepsilon/4(|I|+1)^2$ -close to an event in $\mathcal{B}_{\alpha,i}$. Since there are only finitely many i , we can make this α uniform in i . This implies in particular that $\|\mathbf{I}(E_i) - \mathbf{P}(E_i|\mathcal{B}_{\alpha,i})\|_{L^2} \leq \varepsilon^{1/2}/2(|I|+1)$ for all $i \in I$, since the orthogonal projection $\mathbf{P}(E_i|\mathcal{B}_{\alpha,i})$ is the nearest $\mathcal{B}_{\alpha,i}$ -measurable random variable to $\mathbf{I}(E_i)$ in the L^2 metric.

Let $E_{\alpha,i} \in \mathcal{B}_{\alpha,i}$ denote the event that $\mathbf{P}(E_i|\mathcal{B}_{\alpha,i}) > \frac{|I|}{|I|+1}$ (this event is only defined up to null events in $\mathcal{B}_{\alpha,i}$). Then by Chebyshev's inequality we have

$$\begin{aligned} \mathbf{P}(E_i \setminus E_{\alpha,i}) &\leq \mathbf{P}\left(|\mathbf{I}(E_i) - \mathbf{P}(E_i|\mathcal{B}_{\alpha,i})| > \frac{1}{|I|+1}\right) \\ &\leq \varepsilon/2 \end{aligned}$$

for each $i \in I$. Now let A denote the event $A := \bigwedge_{i \in I} E_{\alpha,i}$. Since $\mathbf{P}(\bigwedge_{i \in I} E_i) = 0$, we have

$$\mathbf{P}(A) \leq \sum_{i \in I} \mathbf{P}(A \setminus E_i). \quad (14)$$

On the other hand, we have

$$\mathbf{P}(A \setminus E_i) = \mathbf{E}(\mathbf{I}(E_{\alpha,i} \setminus E_i) \prod_{j \in I \setminus \{i\}} \mathbf{I}(E_{\alpha,j})).$$

Using the relative independence hypothesis, we conclude

$$\mathbf{P}(A \setminus E_i) = \mathbf{E}(\mathbf{P}(E_{\alpha,i} \setminus E_i | \mathcal{B}_{\alpha,i}) \prod_{j \in I \setminus \{i\}} \mathbf{I}(E_{\alpha,j})).$$

But by definition of $E_{\alpha,i}$ we have

$$\mathbf{P}(E_{\alpha,i} \setminus E_i | \mathcal{B}_{\alpha,i}) = \mathbf{I}(E_{\alpha,i})(1 - \mathbf{P}(E_i | \mathcal{B}_{\alpha,i})) \leq \frac{1}{|I| + 1} \mathbf{I}(E_{\alpha,i})$$

and hence

$$\mathbf{P}(A \setminus E_i) \leq \frac{1}{|I| + 1} \mathbf{E}(\mathbf{I}(E_{\alpha,i}) \prod_{j \in I \setminus \{i\}} \mathbf{I}(E_{\alpha,j})) = \frac{1}{k + 1} \mathbf{P}(A).$$

Inserting this back into (14) we conclude that $\mathbf{P}(A) \leq \frac{|I|}{|I| + 1} \mathbf{P}(A)$, and hence that A is a null event. By definition of A and the UIP hypothesis, we may thus find regular events $F_{\alpha,i} \in \mathcal{B}_{\alpha,i}$ for all $i \in I$ with $\mathbf{P}(E_{\alpha,i} \setminus F_{\alpha,i}) \leq \varepsilon/2$ such that $\bigwedge_{i \in I} F_{\alpha,i} = \emptyset$. By the triangle inequality we thus have $\mathbf{P}(E_i \setminus F_{\alpha,i}) \leq \varepsilon/2$. The claim now follows by setting $F_i := F_{\alpha,i}$. \blacksquare

6. PROOF OF THE INFINITARY HYPERGRAPH REMOVAL LEMMA

We are now ready to prove the hypergraph removal lemma. Fix the probability space $(\Omega, \mathcal{B}_{\max}, \mathbf{P})$, the algebra \mathcal{B}_{reg} of regular events, the finite set, J , the downset \mathbf{i}_{\max} , and the factors \mathcal{B}_I for $I \in \mathbf{i}_{\max}$ obeying the hypotheses in Theorem 4.2. For any sub-ideal \mathbf{i} of \mathbf{i}_{\max} , let $\mathcal{B}(\mathbf{i})$ denote the factor $\mathcal{B}(\mathbf{i}) := \bigvee_{I \in \mathbf{i}} \mathcal{B}_I$, thus $\mathcal{B}(\mathbf{i})$ is a regularisable factor. For any $I \in \mathbf{i}_{\max}$, define the *principal ideal* $\langle I \rangle := \{I' : I' \subseteq I\}$; from the nesting property we see that $\mathcal{B}(\langle I \rangle) = \mathcal{B}_I$ for all $I \in \mathbf{i}_{\max}$. Thus our task is to show that the tuple $(\mathcal{B}(\langle I \rangle))_{I \in \mathbf{i}_{\max}}$ obeys the UIP. For inductive purposes, we will derive this claim from the following more general statement. For any downset \mathbf{i} , we define the *height* $h(\mathbf{i})$ of \mathbf{i} to be the quantity $h(\mathbf{i}) := \sup\{|e| : e \in \mathbf{i}\}$, with the convention that the empty ideal has height $-\infty$.

Proposition 6.1. *Let the hypotheses and notation be as above. Let $d \geq 0$, and let $(\mathbf{i}_i)_{i \in I}$ be any finite tuple of sub-ideals of \mathbf{i}_{\max} (possibly with repetitions), such that every ideal \mathbf{i}_i has height at most d . Then the tuple $(\mathcal{B}(\mathbf{i}_i))_{i \in I}$ obeys the UIP.*

By taking d sufficiently large (e.g. $d = |J|$) we obtain Theorem 4.2.

Proof We will prove Proposition 6.1 by an induction on d . First consider the base case $d = 0$. Then the only ideals available are the empty ideal $\{\}$, and the singleton ideal $\{\emptyset\}$; these correspond to the trivial factor $\{\emptyset, \Omega\}$ and the regularisable factor \mathcal{B}_{\emptyset} . The claim now follows from Examples 5.3, 5.4, 5.6 and Lemma 5.11.

Now suppose that $d \geq 1$, and that Proposition 6.1 has already been proven for $d - 1$. First observe from Lemma 5.11 that we may remove duplicates and assume that all the ideals \mathbf{i}_i are distinct.

Given any $e \in \mathbf{i}_{\max}$ with $|e| = d$, we know that \mathcal{B}_e is regularisable, hence we may write $\mathcal{B}_e = \bigvee_{n \geq 1} \mathcal{B}_{e,n}$ for some increasing sequence $\mathcal{B}_{e,1} \subseteq \mathcal{B}_{e,2} \subseteq \dots$ of

regularisable finite σ -algebras. In particular, we have $\mathcal{B}(\mathbf{i}_i) = \bigvee_{n \geq 1} \mathcal{B}_n(\mathbf{i}_i)$ for all $i \in I$, where

$$\mathcal{B}_n(\mathbf{i}) := \bigvee_{e \in \mathbf{i}: |e|=d} \mathcal{B}_{e,n} \vee \mathcal{B}(\bar{\mathbf{i}})$$

and $\bar{\mathbf{i}}$ is the downset $\bar{\mathbf{i}} := \{e \in \mathbf{i} : |e'| < d\}$; note that this ideal has height strictly less than d .

We need some relative independence properties of the factors $\mathcal{B}(\mathbf{i})$. We begin with

Lemma 6.2. *Let \mathbf{i}, \mathbf{i}' be sub-ideals of \mathbf{i}_{\max} height at most d which do not have any common elements of order exactly d . Then $\mathcal{B}(\mathbf{i})$ and $\mathcal{B}(\mathbf{i}')$ are relatively independent over $\mathcal{B}(\bar{\mathbf{i}})$.*

Proof We will induct on the quantity $m := |\{e \in \mathbf{i} : |e| = d\}|$, the number of top-order elements in \mathbf{i} . If $m = 0$ then $\mathbf{i} = \bar{\mathbf{i}}$ and the claim follows. Now suppose that $m \geq 1$ and the claim has already been established for $m - 1$. Let e_d be an element of \mathbf{i} with $|e_d| = d$, and let $\tilde{\mathbf{i}} := \mathbf{i} \setminus \{e_d\}$. From the induction hypothesis we already know that $\mathcal{B}(\tilde{\mathbf{i}})$ and $\mathcal{B}(\mathbf{i}')$ are relatively independent over $\mathcal{B}(\bar{\mathbf{i}})$. Also, from the UIP hypothesis we know that $\mathcal{B}(\langle e_d \rangle)$ and $\mathcal{B}(\tilde{\mathbf{i}}) \vee \mathcal{B}(\mathbf{i}')$ are relatively independent over $\mathcal{B}(\bar{\mathbf{i}})$. Applying the gluing property (Proposition A.27(i)) we conclude that the factors $\mathcal{B}(\tilde{\mathbf{i}}) \vee \mathcal{B}(\langle e_d \rangle)$ and $\mathcal{B}(\mathbf{i}')$ are relatively independent over $\mathcal{B}(\bar{\mathbf{i}})$. Since the former factor is nothing more than $\mathcal{B}(\mathbf{i})$, the claim follows. \blacksquare

As a consequence, we have

Lemma 6.3. *Let $i \in I$ and $n \geq 1$. Then $\mathcal{B}(\mathbf{i}_i)$ and $\bigvee_{j \in I: j \neq i} \mathcal{B}_n(\mathbf{i}_j)$ are relatively independent over $\mathcal{B}_n(\mathbf{i}_i)$.*

Proof Observe that $\bigvee_{j \in I: j \neq i} \mathcal{B}_n(\mathbf{i}_j)$ is a factor of $\mathcal{B}_n(\mathbf{i}_i) \vee \mathcal{B}(\mathbf{i}')$, where \mathbf{i}' is the downset $\mathbf{i}' := (\bigcup_{j \in I} \mathbf{i}_j) \setminus \{e \in \mathbf{i}_i : |e| = d\}$. Thus by monotonicity and absorption (Proposition A.27(i), (ii)) it suffices to show that $\mathcal{B}(\mathbf{i}_i)$ and $\mathcal{B}(\mathbf{i}')$ are relatively independent over $\mathcal{B}_n(\mathbf{i}_i)$. Since factors do not affect relative independence (Proposition A.27(iv)), it suffices to show that $\mathcal{B}(\mathbf{i}_i)$ and $\mathcal{B}(\mathbf{i}')$ are relatively independent over $\mathcal{B}(\bar{\mathbf{i}}_i)$. But this follows from Lemma 6.2. \blacksquare

From the above lemma and Lemma 5.14, we see that to close the induction hypothesis it suffices to show that $(\mathcal{B}_n(\mathbf{i}_i))_{i \in I}$ obeys the UIP for all $n \geq 1$.

Let k denote the number of ideals \mathbf{i}_i whose height is exactly d . First suppose that all the ideals \mathbf{i}_i have height strictly less than d . Then $\mathcal{B}_n(\mathbf{i}_i) = \mathcal{B}(\mathbf{i}_i)$, and the claim follows from the induction hypothesis.

Now suppose that all the ideals \mathbf{i}_i either have height strictly less than d , or are principal ideals (this is a “weakly mixing” case). We induct on the number of principal ideals of height d . If there are no such ideals, then we are done by the preceding paragraph. Since we have removed duplicates, we know that no two principal ideals present have any common elements of top order d . Thus if \mathbf{i}_i is a

principal ideal, then $\mathcal{B}_n(\mathbf{i}_i)$ is relatively independent of $\bigvee_{j \in I: j \neq i} \mathcal{B}_n(\mathbf{i}_j)$ over $\mathcal{B}(\overline{\mathbf{i}_i})$. Applying Lemma 5.12, it suffices for the purposes of checking the UIP to replace $\mathcal{B}_n(\mathbf{i}_i)$ with $\mathcal{B}_n(\overline{\mathbf{i}_i})$. But this follows from the (inner) induction hypothesis.

Finally, we consider the general case. Let k denote the number of ideals \mathbf{i}_i of height d which are not principal. We have already dealt with the case $k = 0$, so suppose inductively that $k \geq 1$ and the claim has already been proven for $k - 1$. Let \mathbf{i}_{i_0} be an ideal of height d which is not principal, and let e_1, \dots, e_l be the elements of \mathbf{i}_{i_0} of order d . We can then split

$$\mathcal{B}_n(\mathbf{i}_{i_0}) = \mathcal{B}_n(\langle e_1 \rangle) \vee \dots \vee \mathcal{B}_n(\langle e_l \rangle) \vee \mathcal{B}(\overline{\mathbf{i}_i}).$$

Also observe that for $1 \leq j \leq l$ we have $\mathcal{B}_n(\langle e_j \rangle) = \mathcal{B}_{e_j, n} \vee \mathcal{B}(\overline{\langle e_j \rangle})$, and that $\mathcal{B}(\overline{\langle e_j \rangle})$ is a factor of $\mathcal{B}(\overline{\mathbf{i}_i})$. Thus we may apply Lemma 5.13 and conclude that in order to prove the UIP for $(\mathcal{B}_n(\mathbf{i}_i))_{i \in I}$, it suffices to do so for the tuple

$$(\mathcal{B}_n(\mathbf{i}_i))_{i \in I \setminus \{i_0\}} \uplus (\mathcal{B}_n(\langle e_1 \rangle), \dots, \mathcal{B}_n(\langle e_l \rangle), \mathcal{B}(\overline{\mathbf{i}_i})).$$

This tuple has one fewer non-principal degree d ideal than the original tuple, and so the claim now follows from the (inner) induction hypothesis. \blacksquare

7. A HYPERGRAPH CORRESPONDENCE PRINCIPLE

We now generalise the graph correspondence principle developed in Section 3 to hypergraphs. To keep the exposition somewhat simple we shall restrict our attention to the principle for a single d -uniform hypergraphs, although there would be no difficulty extending this principle to systems of hypergraphs of varying uniformities and partite-ness. The material here will be extremely analogous to Section 3. Indeed, we could have deleted that section as being redundant, but we believe for pedagogical purposes that it is better to start with graphs before moving on to hypergraphs.

Definition 7.1 (Hypergraphs). Let $d \geq 0$. If V is a set, we let $\binom{V}{d} := \{e \subset V : |e| = d\}$ denote the d -element subsets of V . A d -uniform hypergraph is a pair $G = (V, E)$, where V is a non-empty set and $E \subset \binom{V}{d}$.

Note that a 2-uniform hypergraph is the same concept as an undirected graph. We will fix $d \geq 2$, and consider the question of extracting an infinitary limit from a sequence $G^{(m)} = (V^{(m)}, E^{(m)})$ of d -uniform hypergraphs. As before, we shall need a universal space, an embedding into that space, and a correspondence principle. We begin with the universal space.

Definition 7.2 (Hypergraph universal space). Fix $d \geq 2$. Let $\Omega := 2^{\binom{\mathbb{N}}{d}} = \{(\mathbb{N}, E_\infty) : E_\infty \subset \binom{\mathbb{N}}{d}\}$ denote the space of all infinite d -uniform hypergraphs (\mathbb{N}, E_∞) on the natural numbers. On this space Ω , we introduce the events A_e for all $e \in \binom{\mathbb{N}}{d}$ by $A_e := \{(\mathbb{N}, E_\infty) \in \Omega : e \in E_\infty\}$, and let \mathcal{B}_{\max} be the σ -algebra generated by the A_e . We also introduce the regular algebra \mathcal{B}_{reg} generated by the A_e , thus these are the events that depend only finitely many of the A_e . For any $\sigma \in S_\infty$, we define the associated action on \mathcal{B}_{\max} by mapping $\sigma : A_e \mapsto A_{\sigma(e)}$

and extending this to a σ -algebra isomorphism in the unique manner. For any (possibly infinite) subset I of \mathbf{N} , we define \mathcal{B}_I to be the factor of \mathcal{B}_{\max} generated by the events A_e for $e \in \binom{I}{d}$.

Next, we need a way to embed every finite hypergraph into the universal space.

Definition 7.3 (Hypergraph universal embedding). Fix $d \geq 2$. Let $m \geq 1$, and let $G^{(m)} = (V^{(m)}, E^{(m)})$ be a finite d -uniform hypergraph. Let $(\Omega^{(m)}, \mathcal{B}_{\max}^{(m)}, \mathbf{P}^{(m)})$ be the probability space corresponding to the sampling of a countable sequence of iid random variables $x_1^{(m)}, x_2^{(m)}, \dots \in V^{(m)}$ sampled independently and uniformly at random. To every sequence $(x_1^{(m)}, x_2^{(m)}, \dots) \in \Omega^{(m)}$ we associate an infinite d -uniform hypergraph $G_\infty^{(m)} = (\mathbf{N}, E_\infty^{(m)}) \in \Omega$ by setting

$$E_\infty^{(m)} := \{e \in \binom{\mathbf{N}}{d} : \{x_i^{(m)} : i \in e\} \in E^{(m)}\}.$$

This mapping from $\Omega^{(m)}$ to Ω is clearly measurable, since the inverse images of the generating events $A_e \in (\Omega, \mathcal{B}_{\max})$ are the events that $\{x_i^{(m)} : i \in e\}$ lie in $G^{(m)}$, which are certainly measurable in $\mathcal{B}_{\max}^{(m)}$. This allows us to extend the probability measure $\mathbf{P}^{(m)}$ from $(\mathcal{B}_{\max}^{(m)}, \Omega^{(m)})$ to the product space $(\mathcal{B}_{\max} \times \mathcal{B}_{\max}^{(m)}, \Omega \times \Omega^{(m)})$ in a canonical manner, identifying the events A_e with the events $\{x_i^{(m)} : i \in e\} \in E^{(m)}$. We shall abuse notation and refer to the extended measure also as $\mathbf{P}^{(m)}$.

As before we can verify the permutation invariance (4). By repeating the proof of the graph correspondence principle (Proposition 3.4) almost word-for-word, we obtain its counterpart for hypergraphs:

Theorem 7.4 (Hypergraph correspondence principle). Fix $d \geq 2$. For every $m \geq 1$, let $G^{(m)} = (V^{(m)}, E^{(m)})$ be a finite d -uniform hypergraph, and let $\mathbf{P}^{(m)}$ be as in Definition 7.3. Then there exists a subsequence $0 < m_1 < m_2 < \dots$ of m , and a probability measure $\mathbf{P}^{(\infty)}$ on the hypergraph universal space $(\Omega, \mathcal{B}_{\max})$, such that we have the weak convergence property (1) and the permutation invariance property (5). Furthermore, we have the following relative independence property: for any $I, I_1, \dots, I_l \in \mathbf{N}$ with $I \cap I_1 \cap \dots \cap I_l$ infinite, the factors \mathcal{B}_I and $\bigvee_{i=1}^l \mathcal{B}_{I_i}$ are relatively independent conditioning on $\bigvee_{i=1}^l \mathcal{B}_{I \cap I_i}$, with respect to this probability measure $\mathbf{P}^{(\infty)}$.

Similarly, by repeating the proof of Lemma 3.5 almost word for word we obtain

Lemma 7.5 (Infinitary hypergraph regularity lemma). Fix $d \geq 2$, and let \mathbf{P} be a probability measure on the hypergraph universal space $(\Omega, \mathcal{B}_{\max})$ which is permutation-invariant in the sense of (5). Then for any $I, I_1, \dots, I_l \in \mathbf{N}$ with $I \cap I_1 \cap \dots \cap I_l$ infinite, the factors \mathcal{B}_I and $\bigvee_{i=1}^l \mathcal{B}_{I_i}$ are relatively independent conditioning on $\bigvee_{i=1}^l \mathcal{B}_{I \cap I_i}$.

Remark 7.6. We should emphasise just how easily the regularity lemma has extended to the hypergraph case here. This is contrast with the development of the finitary hypergraph regularity lemma, which has only been satisfactorily achieved

quite recently [19], [20], [12], [30] (with preliminary work in [5], [3], [6]). In the author's view this is because the regularity lemma is a relatively "soft" component of the theory; in the infinitary framework, the "hard" components of the theory are now isolated in the three fundamental extension properties in Lemma 5.12, Lemma 5.13, Lemma 5.14 (and to a lesser extent in Lemma 5.11). These three lemmas are roughly analogous to the "counting lemma" components of the hypergraph theory (although Lemma 5.14 also captures some of the nature of the "regularity lemma" component, and is the step which is most responsible for the extremely poor quantitative bounds in this theory). Unsurprisingly, it is also these three lemmas where one does the most non-trivial manipulation of small quantities such as ε . Fortunately, the infinitary setting allows one to isolate these epsilons from each other, despite the fact that all three of these basic lemmas are used repeatedly in the proof of the infinitary hypergraph removal lemma (Theorem 4.2). If instead we expanded out all of these lemmas within the proof of Theorem 4.2, and allowed the various epsilons to mix together (with the order of quantifiers, etc. being carefully recorded), one would eventually end up with a complicated situation roughly analogous to those in the finitary proofs [19], [20], [22], [23], [12], [30] of the hypergraph removal lemma. Thus the infinitary perspective allows for a powerful *encapsulation* of distinct components of the argument which greatly cleans up and clarifies the high-level structure of the proof, though the low-level components are, at a fundamental level, essentially the same as in the finitary approach.

8. AN INFINITARY PROOF OF THE HYPERGRAPH REMOVAL LEMMA

We can now repeat the arguments from Section 4 to obtain the following triangle-removal lemma of Nagle, Schacht, Rödl, and Skokan [19], [20], [22], [23] (and independently by Gowers [12]; see also [30] for a later proof):

Theorem 8.1 (Hypergraph removal lemma). *Fix $d \geq 2$, and let $G_0 = (V_0, E_0)$ be a d -uniform hypergraph. Let $G = (V, E)$ be a d -uniform hypergraph with $|V| = n$ vertices. Suppose that G contains fewer than $\delta n^{|V_0|}$ copies of G_0 for some $0 < \delta \leq 1$, or more precisely*

$$|\{(x_i)_{i \in V_0} \in V^{V_0} : \{x_i : i \in e\} \in E \text{ for all } e \in E_0\}| \leq \delta n^{|V_0|}.$$

Then it is possible to delete $o_{\delta \rightarrow 0; G_0, d}(n^d)$ edges from G to create a d -uniform hypergraph G' which has no copies of G_0 whatsoever. Here the subscripting of the $o()$ notation by G_0, d indicates that the quantity $o_{\delta \rightarrow 0; G_0}(n^d)$, when divided by n^d , goes to zero as $\delta \rightarrow 0$ for each fixed G_0, d , but the decay rate is not uniform in G_0, d .

Remark 8.2. As with the triangle removal lemma, this lemma has previously only been proven via a hypergraph regularity lemma, followed by a counting lemma. This is rather complicated; the shortest proof known (in [30]) is about 25 pages, and requires some quite delicate computations. While this current proof is arguably longer than the proof in [30], and certainly less elementary, there are far fewer computations involved, and we believe the argument here is more conceptually clear. This theorem has a number of applications, most notably in giving a proof not only of Szemerédi's theorem (Theorem 2.1) but also a multidimensional version due to Furstenberg and Katznelson [9]; see e.g. [23], [12], [31] for further discussion of this connection, and [21] for some more applications of this theorem. A variant

of this theorem was also used in [31] to establish that the Gaussian primes contain arbitrarily shaped constellations; we shall discuss this variant shortly.

Proof (Sketch) This is basically a repetition of the proof of Lemma 4.1, so we sketch the main points only. Fix d, G_0 . We can relabel V_0 to be $\{1, \dots, n_0\}$ for some integer n_0 ; we can also easily assume that E_0 is non-empty. If the theorem failed, we argue much as in the proof of Lemma 4.1, with $\{1, \dots, n_0\}$ playing the role of $\{1, 2, 3\}$ (and thus $\{n_0 + 1, n_0 + 2, \dots\}$ playing the role of $\{4, 5, \dots\}$). We apply the hypergraph correspondence principle to obtain an infinitary limiting system $(\Omega, \mathcal{B}_{\max}, \mathbf{P}^{(\infty)})$, and apply Theorem 4.2 with $J := \{1, \dots, n_0\}$, $\mathbf{i}_{\max} := \{e : e \subseteq e' \text{ for some } e' \in V_0\}$, E_e set equal to A_e if $e \in V_0$ and $E_e = \Omega$ otherwise (the latter happens precisely when $|e| < d$), and with \mathcal{B}_e set equal to $\mathcal{B}_{e \cup \{n_0+1, n_0+2, \dots\}}$ for all $e \in \mathbf{i}_{\max}$. One then continues the argument as in Lemma 4.1 (with the factor 100 in (8) replaced by at least $2^d d! |E_0|$); the remainder of the proof proceeds with only the obvious minor changes. \blacksquare

Remark 8.3. These results have analogues for partite hypergraphs (see [30]) and are proven similarly, but we will not do so here; the main difference is that instead of sampling all vertices from a single vertex class, one samples countably many vertices from each vertex class (which also leads to a more complicated symmetry group than S_∞). Just as the triangle removal lemma, Lemma 4.1, has a stronger version in Lemma 4.4 which gives a complexity bound on the approximating graph G' , the hypergraph removal lemma given above also comes with a stronger version, in which the approximating hypergraph G' is no longer a subhypergraph of G , but can be described using a partition of $\binom{V}{d-1}$ into $O_{\delta, G_0, d}(1)$ components. We will neither state nor prove this stronger version here (the proof is much the same as Lemma 4.4), but see [30] for an extremely similar statement (in the setting of partite hypergraphs rather than non-partite hypergraphs). This version played an important role in the result in [31] that the Gaussian primes contained arbitrarily shaped constellations.

APPENDIX A. REVIEW OF PROBABILITY THEORY

In this appendix we review the notation and tools from probability that we shall need. There are two concepts here of particular importance: the concept of *relative independence* of two or more factors in a probability space; and the ability to approximate complicated events or random variables by combinations of more elementary events or random variables.

A.1. The algebra of events. A probability space has two major structures; the set-theoretic structure of its events, and the measure-theoretic structure of the probability measure \mathbf{P} . Because we will be dealing with multiple event spaces with a single probability measure, or multiple probability measures on a single event space, it will be conceptually clearer if we treat these two structures separately. We begin with the structure of the event spaces. For technical reasons it is convenient to restrict attention to countably generated spaces.

Definition A.2 (Event spaces). An *event space* is a pair $(\Omega, \mathcal{B}_{\max})$, where the *sample space* Ω is a non-empty set (possibly infinite), and \mathcal{B}_{\max} is a σ -algebra on Ω , i.e. a collection of subsets of Ω which are closed under countable unions, intersections, and complements, and which contains the empty set and Ω . We will also require that the σ -algebra \mathcal{B}_{\max} be *countably generated*, thus there exists a countable sequence of events $E_1, E_2, \dots \in \mathcal{B}_{\max}$ such that \mathcal{B}_{\max} is the minimal σ -algebra containing all these events. We refer to elements of \mathcal{B}_{\max} as (*measurable*) *events*; we abuse notation and identify properties $P(x)$ of points $x \in \Omega$ with the associated event $\{x \in \Omega : P(x) \text{ is true}\}$, and refer to the event simply as P . If A and B are events, we use $A \vee B$ to denote the event that at least one of A and B are true (i.e. $A \vee B$ is the union of A and B) and $A \wedge B$ to denote the event that A and B are not true (i.e. $A \wedge B$ is the intersection of A and B). We also use \overline{A} to denote the event that A is not true (thus $\overline{A} = \Omega \setminus A$).

Example A.3. If Ω is at most countable, the *power-set event space* $(\Omega, 2^\Omega)$ of a set Ω is achieved by setting $\mathcal{B}_{\max} := 2^\Omega := \{E : E \subseteq \Omega\}$ to be the power set of Ω . (If Ω is uncountable, 2^Ω is no longer countably generated.)

Definition A.4 (Factors). Let $(\Omega, \mathcal{B}_{\max})$ be an event space. A *factor* is a subset \mathcal{B} of \mathcal{B}_{\max} which is also a countably generated σ -algebra. More generally, we say that \mathcal{B}_1 is a *factor of* \mathcal{B}_2 (or \mathcal{B}_2 *extends* \mathcal{B}_1) if $\mathcal{B}_1, \mathcal{B}_2$ are both σ -algebras in \mathcal{B}_{\max} and $\mathcal{B}_1 \subseteq \mathcal{B}_2$. We say that a factor is *finite* if it consists of only finitely many events, thus for instance the *trivial factor* $\{\emptyset, \Omega\}$ is finite. An event is \mathcal{B} -*measurable* if it lies in \mathcal{B} . A *random variable* is any function $f : \Omega \rightarrow \mathbf{R}$ with the property that the events $f \in V$ are \mathcal{B}_{\max} -measurable for all open sets V ; if these events are in fact \mathcal{B} -measurable, we say that the random variable f is \mathcal{B} -measurable also. In particular, if an event E is \mathcal{B} -measurable, then its *indicator variable* $\mathbf{I}(E)$, defined to equal 1 when E is true and 0 otherwise, is also \mathcal{B} -measurable. If $\mathcal{E} \subseteq \mathcal{B}_{\max}$ is any collection of events, we let $\mathcal{B}[\mathcal{E}]$ denote the factor *generated by* these events (i.e. the intersection of all factors that contain \mathcal{E}). In particular, if E is a single event, we let $\mathcal{B}[E] = \{\emptyset, E, \overline{E}, \Omega\}$ denote the (finite) factor generated by E . Similarly, if X is a random variable taking finitely many values, we use $\mathcal{B}[X]$ to denote the factor generated by the events $X = c$, where c ranges over the range of X . We write $\mathcal{B}_1 \vee \mathcal{B}_2$ for $\mathcal{B}[\mathcal{B}_1 \cup \mathcal{B}_2]$, thus $\mathcal{B}_1 \vee \mathcal{B}_2$ is the least common extension of \mathcal{B}_1 and \mathcal{B}_2 . More generally, we can define the least common extension $\bigvee_{\alpha \in A} \mathcal{B}_\alpha$ of any at most countable collection of factors \mathcal{B}_α .

Example A.5 (Finite factors). Let A_1, \dots, A_n be a partition of the sample space Ω into disjoint non-empty events. Then $\mathcal{B} = \mathcal{B}[A_1, \dots, A_n]$ is the finite factor consisting of all events which are the union of zero or more of the A_i (and all finite factors are of this form). We refer to A_i as the *atoms* of \mathcal{B} . Let $i : \Omega \rightarrow \{1, \dots, n\}$ be the random variable which indexes which atom one lies in, thus $x \in A_{i(x)}$ for all $x \in \Omega$. A random variable f is \mathcal{B} -measurable if and only if it is determined by i , thus $f(x) = F(i(x))$ for some function $F : \{1, \dots, n\} \rightarrow \mathbf{R}$. One finite factor \mathcal{B}_1 extends another \mathcal{B}_2 if the partition into \mathcal{B}_1 -atoms is finer than the partition into \mathcal{B}_2 -atoms (thus every \mathcal{B}_2 -atom is the union of \mathcal{B}_1 -atoms).

We shall also need the notion of a (*boolean*) *algebra*, namely a subset \mathcal{B} of \mathcal{B}_{\max} which is closed under *finite* intersections, unions, complements, and contains \emptyset and Ω . Thus every factor is an algebra, but not conversely. The reason we need

to deal with algebras rather than factors is because of the observation that the algebra generated by a countable sequence of events remains countable (indeed it is nothing more than the collection of finite boolean combinations of events from that sequence), whereas the factor generated by the same sequence can be uncountable. This is important when applying the Arzela-Ascoli diagonalisation argument (see Lemma A.15 below).

Example A.6. Let $\Omega = [0, 1]^2$, and let \mathcal{B}_{\max} be the Borel σ -algebra (i.e. the algebra generated by the open sets). Let \mathcal{B}_{reg} be the space of *elementary sets*, defined as the finite unions of half-open rectangles $[a, b) \times [c, d)$ where a, b, c, d are rational. Then \mathcal{B}_{reg} is an algebra but not a factor, and is countable; furthermore \mathcal{B}_{\max} is generated by \mathcal{B}_{reg} .

A.7. Probability spaces. We now add the structure of a probability measure to an event space, to form a probability space.

Definition A.8 (Probability spaces). A *probability space* is a triplet $(\Omega, \mathcal{B}_{\max}, \mathbf{P})$, where $(\Omega, \mathcal{B}_{\max})$ is an event space, and $\mathbf{P} : \mathcal{B}_{\max} \rightarrow [0, 1]$ is a *probability measure*, i.e. a countably additive non-negative measure on \mathcal{B}_{\max} with $\mathbf{P}(\Omega) = 1$. A *null event* is an event of probability zero. A statement is true *almost surely* if it is only false on a null event.

Remark A.9. We do not assume our event space $(\Omega, \mathcal{B}_{\max})$ to be complete. Thus, it is not necessarily the case that any subset of a null event is still a measurable event. (It may help to think of the σ -algebras here as being like Borel σ -algebras — that is, algebras generated by open sets — rather than Lebesgue σ -algebras.)

In the remainder of this appendix we assume that the probability space $(\Omega, \mathcal{B}_{\max}, \mathbf{P})$ is fixed.

Definition A.10 (Random variables). We consider two random variables equivalent if they are almost surely equal. If f is absolutely integrable, we use $\mathbf{E}(f)$ to denote the integral of f with respect to the probability measure \mathbf{P} , and write $\|f\|_{L^1} = \|f\|_{L^1(\mathcal{B}_{\max}; \mathbf{P})}$ for $\mathbf{E}(|f|)$. Thus for instance $\mathbf{E}(\mathbf{I}(E)) = \mathbf{P}(E)$ for any event E . Similarly, we write $\|f\|_{L^2} = \|f\|_{L^2(\mathcal{B}_{\max}; \mathbf{P})}$ for $\mathbf{E}(|f|^2)^{1/2}$ whenever f is square-integrable, and $\|f\|_{L^\infty} = \|f\|_{L^\infty(\mathcal{B}_{\max}; \mathbf{P})}$ for the essential supremum of f . We will drop the measure \mathbf{P} , and sometimes the factor \mathcal{B}_{\max} , from the $L^p(\mathcal{B}_{\max}; \mathbf{P})$ notation when these are clear from context.

It will be important to develop relative versions of all these concepts with respect to factors of \mathcal{B}_{\max} .

Definition A.11 (Conditional expectation). If $p = 1, 2, \infty$ and \mathcal{B} is a factor, we let $L^p(\mathcal{B}) = L^p(\mathcal{B}; \mathbf{P})$ denote the space of \mathcal{B} -measurable random variables with finite L^p norm (identifying two random variables if they are equivalent). Observe that $L^2(\mathcal{B})$ is a Hilbert space with inner product $\langle f, g \rangle := \mathbf{E}(fg)$; since \mathcal{B} is countably generated, we see that $L^2(\mathcal{B})$ is separable. We define the *conditional expectation operator* $f \mapsto \mathbf{E}(f|\mathcal{B})$ to be the orthogonal projection from $L^2(\mathcal{B}_{\max})$ to $L^2(\mathcal{B})$; note that $\mathbf{E}(f|\mathcal{B})$ is only defined up to almost sure equivalence. If E is an event, we write $\mathbf{P}(E|\mathcal{B})$ for $\mathbf{E}(\mathbf{I}(E)|\mathcal{B})$, and refer to $\mathbf{P}(E|\mathcal{B})$ as the *conditional probability* of E with respect to the factor \mathcal{B} .

We have the useful

Lemma A.12 (Pythagoras' theorem). *Let \mathcal{B}' be an extension of \mathcal{B} . Then for any $f \in L^2(\mathcal{B}_{\max})$ we have*

$$\|\mathbf{E}(f|\mathcal{B}')\|_{L^2}^2 = \|\mathbf{E}(f|\mathcal{B})\|_{L^2}^2 + \|\mathbf{E}(f|\mathcal{B}') - \mathbf{E}(f|\mathcal{B})\|_{L^2}^2.$$

Proof This follows since $\mathbf{E}(f|\mathcal{B}')$ is the orthogonal projection to $L^2(\mathcal{B}')$, and $\mathbf{E}(f|\mathcal{B})$ is the orthogonal projection to the smaller space $L^2(\mathcal{B})$. ■

Remark A.13. In this paper we shall deal almost exclusively with bounded random variables (indeed, they will almost always be bounded between -1 and 1). Thus issues of *integrability* will not be a concern to us; this also means that we do not have to distinguish between convergence in L^1 , convergence in L^2 , and convergence in measure. It will however be crucial to keep track the *measurability* of our random variables with respect to the various factors involved in the argument.

Example A.14 (Finite factors). Let \mathcal{B} be a finite factor with atoms A_1, \dots, A_n . If $f \in L^\infty(\mathcal{B}_{\max})$, the conditional expectation $\mathbf{E}(f|\mathcal{B})$ is well-defined on all atoms A_i of non-zero probability, and is equal to $\mathbf{E}(f|A_i) := \mathbf{E}(f\mathbf{I}(A_i))/\mathbf{P}(A_i)$ on each such atom. Similarly we have $\mathbf{P}(E|\mathcal{B}) = \mathbf{P}(E|A_i) := \mathbf{P}(E \wedge A_i)/\mathbf{P}(A_i)$ on such atoms. Of course one can develop similar explicit formulae for the conditional covariance of two random variables or events.

We recall some very standard properties of conditional expectation, that we shall use without further comment. The conditional expectation operation $f \mapsto \mathbf{E}(f|\mathcal{B})$ is linear, positivity preserving, and is a contraction on L^p for $p = 1, 2, \infty$. In particular conditional expectation is continuous in each of the L^p topologies, which allows us to easily apply density arguments when verifying identities involving conditional expectation (i.e. it suffices to verify such identities for a dense subclass of random variables, such as simple random variables). We also have the module property that $\mathbf{E}(fg|\mathcal{B}) = f\mathbf{E}(g|\mathcal{B})$ whenever $f \in L^\infty(\mathcal{B})$ and $g \in L^\infty(\mathcal{B}_{\max})$.

In order to pass from a sequence of finitary objects to an infinitary one, the following lemma will be crucial.

Lemma A.15 (Arzela-Ascoli diagonalisation argument). *Let $\mathbf{P}^{(1)}, \mathbf{P}^{(2)}, \dots$ be a sequence of probability measures on an event space $(\Omega, \mathcal{B}_{\max})$. Let \mathcal{B}_{reg} be a countable algebra which generates \mathcal{B}_{\max} as a σ -algebra. Then there exists a subsequence $0 < k_1 < k_2 < \dots$ of integers and a probability measure \mathbf{P} such that*

$$\lim_{i \rightarrow \infty} \mathbf{P}^{(k_i)}(F) = \mathbf{P}(F) \text{ for all } F \in \mathcal{B}_{\text{reg}}.$$

In other words, $\mathbf{P}^{(k_i)}$ is weakly convergent to \mathbf{P} , when tested against the algebra of events \mathcal{B}_{reg} .

Proof We enumerate \mathcal{B}_{reg} as F_1, F_2, \dots (duplicating events if necessary, if \mathcal{B}_{reg} happens to be finite). By using the sequential compactness of the unit interval $[0, 1]$ (i.e. the Heine-Borel theorem), we can obtain a sequence $k_{1,1} < k_{1,2} < k_{1,3} < \dots$ such that $\mathbf{P}^{(k_{1,i})}(F_1)$ converges as $i \rightarrow \infty$ to a limit, say $p_1 \in [0, 1]$. Then we can

extract a subsequence $k_{2,1} < k_{2,2} < k_{2,3} < \dots$ of that sequence such that $\mathbf{P}^{(k_{2,i})}(F_2)$ converges as $i \rightarrow \infty$ to a limit, say $p_2 \in [0, 1]$. We continue in this fashion and then extract the diagonal sequence $k_i := k_{i,i}$ to obtain a sequence $p_1, p_2, \dots \in [0, 1]$ such that $\lim_{i \rightarrow \infty} \mathbf{P}^{(k_i)}(F_j) = p_j$ for each $j = 1, 2, \dots$. One easily verifies that the map $F_j \mapsto p_j$ is *finitely* additive, non-negative, and maps \emptyset to 0 and Ω to 1. Invoking the Kolmogorov extension theorem (or the Carathéodory extension theorem) we can construct a probability measure \mathbf{P} such that $\mathbf{P}(F_j) = p_j$, and the claim follows. ■

Remark A.16. One can also obtain this lemma from the Banach-Alaoglu theorem and the Riesz representation theorem (though one should take care to distinguish the notions of compactness and sequential compactness). Observe that both the Heine-Borel theorem and the Kolmogorov extension theorem are completely constructive, and so this lemma does not use the axiom of choice. See [32] for further discussion.

A.17. Approximation lemmas. We will frequently need to approximate a random variable or event in a complicated factor by linear, polynomial, or boolean combinations of random variables or events in simpler factors. To do this we shall use some very simple and standard tools, which we collect here for the reader's convenience.

Recall that a random variable is *simple* if it only takes on finitely many values, or equivalently if it is the finite linear combination of indicator functions, or equivalently if it is measurable with respect to a finite factor. The following lemma is standard in measure theory:

Lemma A.18. *Let \mathcal{B} be a factor and $p = 1, 2, \infty$. Then the simple random variables in $L^p(\mathcal{B})$ are dense in $L^p(\mathcal{B})$.*

Because of this, the task of approximating random variables quickly boils down to approximating events. Let us say that two events E, F are ε -close if $\mathbf{P}(E \setminus F) + \mathbf{P}(F \setminus E) \leq \varepsilon$.

Lemma A.19 (Approximation by finite complexity events). *Let $\mathcal{B} = \mathcal{B}[\mathcal{E}]$ be a factor generated by a (possibly infinite) collection \mathcal{E} of events, and let $\varepsilon > 0$. Then every event in \mathcal{B} is ε -close to a finite boolean combination of events from \mathcal{E} . In particular, if \mathcal{B} is generated by an algebra \mathcal{B}_{reg} , then every event in \mathcal{B} is ε -close to an event from \mathcal{B}_{reg} . If $f \in L^1(\mathcal{B})$, then there exists a finite factor \mathcal{B}' of \mathcal{B} generated by finitely many events in \mathcal{E} , such that $\|f - \mathbf{E}(f|\mathcal{B}')\|_{L^1(\mathcal{B})} \leq \varepsilon$.*

Proof Let \mathcal{B}_{reg} be the algebra generated by \mathcal{E} (i.e. the space of finite boolean combinations of events from \mathcal{E}). Let \mathcal{B}_ε denote the collection of events which is ε -close to an element of \mathcal{B}_{reg} . Then one easily verifies that $\bigcap_{\varepsilon > 0} \mathcal{B}_\varepsilon$ is a factor that contains \mathcal{E} , and thus contains \mathcal{B} , and the first and second claims follow. To prove the final claim, first use Lemma A.18 to reduce to the case where f is simple, and then use linearity to reduce to the case when $f = \mathbf{I}(E)$ is an indicator function. By the previous claims, we can find an event $E' \in \mathcal{B}_{\text{reg}}$ which is $\varepsilon/2$ -close to E , thus $\|f - \mathbf{I}(E')\|_{L^1(\mathcal{B})} \leq \varepsilon/2$. This E' lies in some finite factor \mathcal{B}' generated by \mathcal{E} , and thus

on taking conditional expectations in \mathcal{B}' we see that $\|\mathbf{E}(f|\mathcal{B}') - \mathbf{I}(E')\|_{L^1(\mathcal{B})} \leq \varepsilon/2$. The claim now follows from the triangle inequality. ■

Corollary A.20 (Limits of chains). *Let A be a totally ordered set, and for each $\alpha \in A$ let \mathcal{B}_α be a factor of \mathcal{B}_{\max} with the monotonicity property $\mathcal{B}_\alpha \subseteq \mathcal{B}_\beta$ whenever $\alpha \leq \beta$. Let $\mathcal{B} := \bigvee_{\alpha \in A} \mathcal{B}_\alpha$. Then for any $f \in L^2(\mathcal{B})$, the net $\mathbf{E}(f|\mathcal{B}_\alpha)$ converges to f in L^2 norm (thus for every $\varepsilon > 0$ there exists $\beta \in A$ such that $\|f - \mathbf{E}(f|\mathcal{B}_\alpha)\|_{L^2(\mathcal{B})} \leq \varepsilon$ for all $\alpha \geq \beta$).*

Proof Let $\varepsilon > 0$. Applying Lemma A.19 with $\mathcal{E} = \bigcup_{\alpha \in A} \mathcal{B}_\alpha$, we can find a finite factor \mathcal{B}' generated by finitely many events in \mathcal{E} such that $\|f - \mathbf{E}(f|\mathcal{B}')\|_{L^2(\mathcal{B})} \leq \varepsilon$. By monotonicity we see that \mathcal{B}' is a factor of \mathcal{B}_α for some $\alpha \in A$. The claim then follows from Pythagoras' theorem. ■

Corollary A.21 (Approximation by finite factors). *Let $\mathcal{B}_1, \dots, \mathcal{B}_k$ be factors and $\varepsilon > 0$. Then every event in $\mathcal{B}_1 \vee \dots \vee \mathcal{B}_k$ is ε -close to a finite boolean combination of events in $\mathcal{B}_1 \cup \dots \cup \mathcal{B}_k$. Furthermore, given any random variable $f \in L^\infty(\mathcal{B}_1 \vee \dots \vee \mathcal{B}_k)$, there exists finite factors \mathcal{B}'_i of \mathcal{B}_i for $i = 1, \dots, k$ respectively such that $\|f - \mathbf{E}(f|\mathcal{B}'_1 \vee \dots \vee \mathcal{B}'_k)\|_{L^1(\mathcal{B}_1 \vee \dots \vee \mathcal{B}_k)} \leq \varepsilon$.*

Proof The first claim follows from Lemma A.19 by setting $\mathcal{E} := \mathcal{B}_1 \cup \dots \cup \mathcal{B}_k$. To verify the second claim, first use Lemma A.19 to locate a finite factor \mathcal{B}' generated by finitely many elements in $\mathcal{B}_1 \cup \dots \cup \mathcal{B}_k$ such that $\|f - \mathbf{E}(f|\mathcal{B}')\|_{L^1(\mathcal{B}_1 \vee \dots \vee \mathcal{B}_k)} \leq \varepsilon/2$. Now observe that \mathcal{B}' is a factor of $\mathcal{B}'_1 \vee \dots \vee \mathcal{B}'_k$ for some finite factors \mathcal{B}'_i of \mathcal{B}_i for $i = 1, \dots, k$. The claim now follows from the same triangle inequality argument used to prove Lemma A.19. ■

A.22. Relative independence. Now we come to a fundamental notion for us, namely that of (relative) independence of two or more factors.

Definition A.23 (Independence). We say that two factors $\mathcal{B}_1, \mathcal{B}_2$ are *unconditionally independent* if we have

$$\mathbf{E}(f_1 f_2) = \mathbf{E}(f_1) \mathbf{E}(f_2)$$

for all $f_1 \in L^\infty(\mathcal{B}_1)$ and $f_2 \in L^\infty(\mathcal{B}_2)$. More generally, we say that two factors $\mathcal{B}_1, \mathcal{B}_2$ are *relatively independent* conditioning on a third factor \mathcal{B} with respect to the probability measure \mathbf{P} if we have

$$\mathbf{E}(f_1 f_2 | \mathcal{B}) = \mathbf{E}(f_1 | \mathcal{B}) \mathbf{E}(f_2 | \mathcal{B}) \quad (15)$$

almost surely for all $f_1 \in L^\infty(\mathcal{B}_1)$ and $f_2 \in L^\infty(\mathcal{B}_2)$. In many cases, the probability measure \mathbf{P} will be clear from context and we shall omit the phrase “with respect to \mathbf{P} ”. Given an at most countable collection of factors $(\mathcal{B}_\alpha)_{\alpha \in A}$, we say that these factors are *jointly unconditionally independent* (resp. *jointly relatively independent* conditioning on a factor \mathcal{B}) if $\bigvee_{\alpha \in A_1} \mathcal{B}_\alpha$ and $\bigvee_{\alpha \in A_2} \mathcal{B}_\alpha$ are unconditionally independent (resp. relatively independent conditioning on \mathcal{B}) for all disjoint subsets A_1, A_2 of A . We say that a collection of events E_1, E_2, \dots is unconditionally independent (resp. relatively independent conditioning on \mathcal{B}) if their associated factors $\mathcal{B}[E_1], \mathcal{B}[E_2], \dots$ are unconditionally independent (resp. relatively independent conditioning on \mathcal{B}).

Examples A.24. Two factors $\mathcal{B}_1, \mathcal{B}_2$ are unconditionally independent if and only if $\mathbf{P}(E \wedge F) = \mathbf{P}(E)\mathbf{P}(F)$ for all $E \in \mathcal{B}_1$ and $F \in \mathcal{B}_2$. In particular, two events E and F are unconditionally independent if and only if $\mathbf{P}(E \wedge F) = \mathbf{P}(E)\mathbf{P}(F)$. Three factors $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ are jointly unconditionally independent if and only if $\mathbf{P}(E_1 \wedge E_2 \wedge E_3) = \mathbf{P}(E_1)\mathbf{P}(E_2)\mathbf{P}(E_3)$ for all $E \in \mathcal{B}_1$ and $F \in \mathcal{B}_2$. On the other hand, in order for three events E, F, G to be jointly independent it is not quite enough that $\mathbf{P}(E \wedge F \wedge G) = \mathbf{P}(E)\mathbf{P}(F)\mathbf{P}(G)$; one also needs E, F, G to be pairwise independent, thus for instance $\mathbf{P}(E \wedge F) = \mathbf{P}(E)\mathbf{P}(F)$. If $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ are jointly unconditionally independent, then $\mathcal{B}_1 \vee \mathcal{B}_3$ and $\mathcal{B}_2 \vee \mathcal{B}_3$ are conditionally independent over \mathcal{B}_3 , even though they are almost certainly not unconditionally independent. On the other hand, \mathcal{B}_1 and \mathcal{B}_2 are both unconditionally independent, and conditionally independent over \mathcal{B}_3 .

Example A.25. Let x_1, x_2, x_3 be three elements chosen uniformly and independently at random from $\{0, 1\}$. Then the events $x_1 = x_3$ and $x_2 = x_3$ are unconditionally independent, but they are not relatively independent conditioning on the factor $\mathcal{B}[x_1 = x_2]$. Thus we see that unconditional independence is neither stronger nor weaker than relative independence.

Taking expectations in (15) we obtain

$$\mathbf{E}(f_1 f_2) = \mathbf{E}(\mathbf{E}(f_1 | \mathcal{B}) f_2) = \mathbf{E}(f_1 \mathbf{E}(f_2 | \mathcal{B})) = \mathbf{E}(\mathbf{E}(f_1 | \mathcal{B}) \mathbf{E}(f_2 | \mathcal{B})) \quad (16)$$

whenever $\mathcal{B}_1, \mathcal{B}_2$ are relatively independent conditioning on \mathcal{B} , $f_1 \in L^\infty(\mathcal{B}_1)$ and $f_2 \in L^\infty(\mathcal{B}_2)$.

There are several equivalent formulations of relative independence.

Lemma A.26. *Let $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}$ be factors. Then the following are equivalent:*

- (i) \mathcal{B}_1 and \mathcal{B}_2 are relatively independent conditioning on \mathcal{B} .
- (ii) We have $\mathbf{E}(f_1 | \mathcal{B} \vee \mathcal{B}_2) = \mathbf{E}(f_1 | \mathcal{B})$ almost surely for all $f_1 \in L^2(\mathcal{B}_1)$.
- (iii) We have $\|\mathbf{E}(f_1 | \mathcal{B} \vee \mathcal{B}_2)\|_{L^2} = \|\mathbf{E}(f_1 | \mathcal{B})\|_{L^2}$ for all $f_1 \in L^2(\mathcal{B}_1)$.
- (iv) We have $\|\mathbf{P}(E_1 | \mathcal{B} \vee \mathcal{B}_2)\|_{L^2} = \|\mathbf{P}(E_1 | \mathcal{B})\|_{L^2}$ for all $E_1 \in \mathcal{B}_1$.

Proof The equivalence of (ii) and (iii) follows from Lemma A.12. The equivalence of (iii) and (iv) follows from Lemma A.18, linearity, and a standard limiting arguments.

To see that (ii) implies (i), observe for $f_1 \in L^\infty(\mathcal{B}_1)$ and $f_2 \in L^\infty(\mathcal{B}_2)$ that

$$\begin{aligned} \mathbf{E}(f_1 f_2 | \mathcal{B}) &= \mathbf{E}(\mathbf{E}(f_1 f_2 | \mathcal{B} \vee \mathcal{B}_2) | \mathcal{B}) \\ &= \mathbf{E}(\mathbf{E}(f_1 | \mathcal{B} \vee \mathcal{B}_2) f_2 | \mathcal{B}) \\ &= \mathbf{E}(\mathbf{E}(f_1 | \mathcal{B}) f_2 | \mathcal{B}) \\ &= \mathbf{E}(f_1 | \mathcal{B}) \mathbf{E}(f_2 | \mathcal{B}) \end{aligned}$$

where we have used the module property twice.

Finally, we show that (i) implies (iv). We observe from (i) and the module property that

$$\mathbf{E}(f_1 f_2 h | \mathcal{B}) = \mathbf{E}(f_1 | \mathcal{B}) \mathbf{E}(f_2 h | \mathcal{B})$$

whenever $f_1 \in L^\infty(\mathcal{B}_1)$, $f_2 \in L^\infty(\mathcal{B}_2)$, and $h \in L^\infty(\mathcal{B})$. Taking linear combinations and using limiting arguments we conclude that

$$\mathbf{E}(f_1 g | \mathcal{B}) = \mathbf{E}(f_1 | \mathcal{B}) \mathbf{E}(g | \mathcal{B})$$

whenever $g \in L^\infty(\mathcal{B} \vee \mathcal{B}_2)$. We take expectations and obtain

$$\mathbf{E}(f_1 g) = \mathbf{E}(\mathbf{E}(f_1 | \mathcal{B}) \mathbf{E}(g | \mathcal{B})).$$

Applying this with $f_1 := \mathbf{I}(E_1)$ and $g = \mathbf{P}(E_1 | \mathcal{B} \vee \mathcal{B}_2)$ we obtain

$$\|\mathbf{P}(E_1 | \mathcal{B} \vee \mathcal{B}_2)\|_{L^2}^2 = \mathbf{E}(\mathbf{I}(E_1) \mathbf{P}(E_1 | \mathcal{B} \vee \mathcal{B}_2)) = \mathbf{E}(\mathbf{P}(E_1 | \mathcal{B}) \mathbf{E}(\mathbf{P}(E_1 | \mathcal{B} \vee \mathcal{B}_2) | \mathcal{B})) = \|\mathbf{P}(E_1 | \mathcal{B})\|_{L^2}^2$$

and (iv) follows. ■

Now we can observe the following stability properties concerning relative independence.

Proposition A.27. *Let $\mathcal{B}_1, \mathcal{B}_2$ be two factors which are relatively independent conditioning on another factor \mathcal{B} .*

- (i) *(Monotonicity) If \mathcal{B}'_1 is a factor of \mathcal{B}_1 and \mathcal{B}'_2 is a factor of \mathcal{B}_2 , then \mathcal{B}'_1 and \mathcal{B}'_2 are relatively independent conditioning on \mathcal{B} .*
- (ii) *(Absorption) $\mathcal{B}_1 \vee \mathcal{B}$ and $\mathcal{B}_2 \vee \mathcal{B}$ are relatively independent conditioning on \mathcal{B} .*
- (iii) *(Gluing) Let \mathcal{B}_3 be a σ -algebra which is relatively independent of $\mathcal{B}_1 \vee \mathcal{B}_2$ conditioning on \mathcal{B} . Then \mathcal{B}_1 is relatively independent of $\mathcal{B}_2 \vee \mathcal{B}_3$ conditioning on \mathcal{B} .*
- (iv) *(Factors do not affect relative independence) If \mathcal{B}'_1 is a factor of \mathcal{B}_1 and \mathcal{B}'_2 is a factor of \mathcal{B}_2 , then \mathcal{B}_1 and \mathcal{B}_2 are relatively independent conditioning on $\mathcal{B} \vee \mathcal{B}'_1 \vee \mathcal{B}'_2$.*
- (v) *(Independent information does not affect relative independence) Let \mathcal{B}_3 be a σ -algebra which is independent of $\mathcal{B} \vee \mathcal{B}_1 \vee \mathcal{B}_2$. Then \mathcal{B}_1 is relatively independent of $\mathcal{B}_2 \vee \mathcal{B}_3$ conditioning on \mathcal{B} .*

Proof The claim (i) is trivial. To prove (ii), observe from symmetry and iteration that it suffices to show that $\mathcal{B}_1 \vee \mathcal{B}$ and \mathcal{B}_2 are relatively independent conditioning on \mathcal{B} . But this follows from two applications of Lemma A.26.

To prove (iii), it suffices by Lemma A.26 (and symmetry) to show that

$$\mathbf{E}(h | \mathcal{B} \vee \mathcal{B}_1) = \mathbf{E}(h | \mathcal{B})$$

for all $h \in L^\infty(\mathcal{B}_2 \vee \mathcal{B}_3)$. By density it suffices to show that

$$\mathbf{E}(f_2 f_3 | \mathcal{B} \vee \mathcal{B}_1) = \mathbf{E}(f_2 f_3 | \mathcal{B})$$

for all $f_2 \in L^\infty(\mathcal{B}_2)$ and $f_3 \in L^\infty(\mathcal{B}_3)$. But this follows from the relative independence hypotheses and the module property:

$$\begin{aligned}
\mathbf{E}(f_2 f_3 | \mathcal{B} \vee \mathcal{B}_1) &= \mathbf{E}(\mathbf{E}(f_2 f_3 | \mathcal{B} \vee \mathcal{B}_1 \vee \mathcal{B}_2) | \mathcal{B} \vee \mathcal{B}_1) \\
&= \mathbf{E}(f_2 \mathbf{E}(f_3 | \mathcal{B} \vee \mathcal{B}_1 \vee \mathcal{B}_2) | \mathcal{B} \vee \mathcal{B}_1) \\
&= \mathbf{E}(f_2 \mathbf{E}(f_3 | \mathcal{B}) | \mathcal{B} \vee \mathcal{B}_1) \\
&= \mathbf{E}(f_2 | \mathcal{B} \vee \mathcal{B}_1) \mathbf{E}(f_3 | \mathcal{B}) \\
&= \mathbf{E}(f_2 | \mathcal{B}) \mathbf{E}(f_3 | \mathcal{B}) \\
&= \mathbf{E}(f_2 f_3 | \mathcal{B}).
\end{aligned}$$

Now we prove (iv). By symmetry and iteration it will suffice to show that \mathcal{B}_1 and \mathcal{B}_2 are relatively independent conditioning on $\mathcal{B} \vee \mathcal{B}'_1$. From Lemma A.26 we already have

$$\|E(f_2 | \mathcal{B})\|_{L^2} = \|\mathbf{E}(f_2 | \mathcal{B} \vee \mathcal{B}_1)\|_{L^2}$$

for all $f_2 \in L^2(\mathcal{B}_2)$. From Lemma A.12 we conclude

$$\|E(f_2 | \mathcal{B})\|_{L^2} = \|\mathbf{E}(f_2 | \mathcal{B} \vee \mathcal{B}'_1)\|_{L^2} = \|\mathbf{E}(f_2 | \mathcal{B} \vee \mathcal{B}_1)\|_{L^2}$$

and the claim follows from another application of Lemma A.26.

Finally, we prove (v). If \mathcal{B}_3 is independent of $\mathcal{B} \vee \mathcal{B}_1 \vee \mathcal{B}_2$, then by the monotonicity and factor properties (i), (iv) we conclude that \mathcal{B}_3 is relatively independent of $\mathcal{B}_1 \vee \mathcal{B}_2$ conditioning on \mathcal{B} . The claim (v) then follows from the gluing property (iii). \blacksquare

APPENDIX B. CONNECTION WITH RECURRENCE THEOREMS

We have just seen how infinitary probabilistic statements such as Theorem 4.2 can imply finitary graph statements such as Lemma 4.1; later we shall see that one can also deduce finitary hypergraph statements in this manner. It is also well known (see [24], [5], [6], [22], [23], [12], [25], [31]) that these graph and hypergraph statements can in turn be used to deduce density results such as Szemerédi's theorem. This in turn is known by the Furstenberg correspondence principle to be equivalent to results such as the Furstenberg recurrence theorem. Concatenating all these implications, one thus expects results such as Theorem 4.2 to be capable of implying results such as Theorem 2.2 directly, without the need to pass back and forth between the infinitary and finitary settings.

Somewhat surprisingly, it appears to be somewhat difficult to achieve this goal; the best the author was able to do was simply to compose the various implications discussed above to obtain a connection. For sake of completeness, we sketch a special case of this connection here, but it is puzzling that there seems to be little “synergy” between these two infinitary results, despite their similarity. As there appear to be no major new features emerging in this connection, we will skip over some of the details.

One can demonstrate the connection using the Furstenberg recurrence theorem (Theorem 2.2), but it will be slightly more convenient to instead work with the following variant:

Theorem B.1 (Furstenberg-Katznelson recurrence theorem, special case). [9] *Let $(\Omega, \mathcal{B}_{\max}, \mathbf{P})$ be a probability space. Let $S, T : \Omega \rightarrow \Omega$ be two commuting probability-preserving bi-measurable maps. Then for all events $A \in \mathcal{B}_{\max}$ with $\mathbf{P}(A) > 0$, we have*

$$\liminf_{N \rightarrow \infty} \frac{1}{2N+1} \sum_{n=-N}^N \mathbf{P}(A \wedge T^n A \wedge S^n A) > 0.$$

This theorem is equivalent to the assertion that any subset of \mathbf{Z}^2 with positive upper density contains infinitely many right-angled triangles $(x, y), (x+r, y), (x, y+r)$, a result first obtained by Ajtai and Szemerédi [1]. In [25] it was observed that this theorem followed from the triangle removal lemma. Setting $S := T^2$ we obtain the $k = 3$ case of Theorem 2.2. The full version of the Furstenberg-Katznelson recurrence theorem allows for an arbitrary number of commuting shifts, and can be treated by a modification of the arguments presented here.

To transfer this theorem to a setting where Theorem 4.2 is applicable, we will have to perform essentially the entire machinery used in the graph correspondence principle. It is convenient not to work with graphs on \mathbf{N} , but rather on tripartite graphs connecting three copies of \mathbf{Z} :

Definition B.2 (Tripartite graph universal space). A *tripartite infinite graph* is a sextuple $G = (\mathbf{Z}, \mathbf{Z}, \mathbf{Z}, E_{12}, E_{23}, E_{31})$ where E_{12}, E_{23}, E_{31} are subsets of \mathbf{Z}^2 . Let Ω^Δ denote the space of all tripartite infinite graphs. On this space Ω^Δ , we introduce the events A_{ij, k_i, k_j} for $ij = 12, 23, 31$ and $k_i, k_j \in \mathbf{Z}$ by $A_{ij, k_i, k_j} := \{G \in \Omega : (k_i, k_j) \in E_{ij}\}$, and let $\mathcal{B}_{\max}^\Delta$ be the σ -algebra generated by the A_{ij, k_i, k_j} . We also introduce the regular algebra $\mathcal{B}_{\text{reg}}^\Delta$ generated by the A_{ij, k_i, k_j} . For any three permutations $\sigma_1, \sigma_2, \sigma_3 : \mathbf{Z} \rightarrow \mathbf{Z}$ we can define an action of $(\sigma_1, \sigma_2, \sigma_3)$ on $\mathcal{B}_{\max}^\Delta$ by mapping A_{ij, k_i, k_j} to $A_{ij, \sigma_i(k_i), \sigma_j(k_j)}$. For any subsets I_1, I_2, I_3 of \mathbf{N} , we define \mathcal{B}_I^Δ to be the factor of $\mathcal{B}_{\max}^\Delta$ generated by the events A_{ij, k_i, k_j} where $ij = 12, 23, 31$, $k_i \in I_i$, and $k_j \in I_j$.

Now we embed the system in Theorem B.1 into this universal space.

Definition B.3 (Tripartite graph universal embedding). Let $(\Omega, \mathcal{B}_{\max}, \mathbf{P})$, S , T , A be as in Theorem B.1. Let $N \geq 1$ be a natural number. We introduce the probability space $(\Omega^{(N)}, \mathcal{B}_{\max}^{(N)}, \mathbf{P}^{(N)})$, defined as the space associated to sampling three infinite sequences $(n_{i, k_i})_{k_i \in \mathbf{Z}}$ for $i = 1, 2, 3$ uniformly and independently at random from $[N] = \{1, \dots, N\}$. Thus the product space $(\Omega \times \Omega^{(N)}, \mathcal{B}_{\max} \times \mathcal{B}_{\max}^{(N)}, \mathbf{P} \times \mathbf{P}^{(N)})$ represents the independent sampling of a point x from Ω , together with the three sequences $n_{i, k_i} \in [N]$ for $i = 1, 2, 3$ and $k_i \in \mathbf{Z}$. For any such x and $n_{i, k}$, we associate an infinite tripartite graph $G = (\mathbf{Z}, \mathbf{Z}, \mathbf{Z}, E_{12}, E_{23}, E_{31})$ in Ω^Δ by setting

$$\begin{aligned} E_{12} &:= \{(k_1, k_2) \in \mathbf{Z} \times \mathbf{Z} : T^{n_{1, k_1}} S^{n_{2, k_2}} x \in A\} \\ E_{23} &:= \{(k_2, k_3) \in \mathbf{Z} \times \mathbf{Z} : T^{n_{3, k_3} - n_{1, k_1}} S^{n_{2, k_2}} x \in A\} \\ E_{31} &:= \{(k_3, k_1) \in \mathbf{Z} \times \mathbf{Z} : T^{n_{1, k_1}} S^{n_{3, k_3} - n_{2, k_2}} x \in A\}. \end{aligned}$$

This is a measurable map from $\Omega \times \Omega^{(N)}$ to Ω^Δ (the inverse image of A_{12,k_1,k_2} is the measurable event $T^{n_1,k_1} S^{n_2,k_2} x \in A$, and similarly for the other two classes of generating events), and so we can push forward the measure $\mathbf{P} \times \mathbf{P}^{(N)}$ to a measure on $(\Omega \times \Omega^{(N)} \times \Omega^\Delta, \mathcal{B}_{\max} \times \mathcal{B}_{\max}^{(N)} \times \mathcal{B}_{\max}^\Delta)$, which by abuse of notation we will also call $\mathbf{P} \times \mathbf{P}^{(N)}$.

A computation (using the probability-preserving and commuting nature of T and S) shows that

$$\begin{aligned} & \mathbf{P} \times \mathbf{P}^{(N)}(A_{12,0,0} \wedge A_{23,0,0} \wedge A_{31,0,0}) \\ &= \frac{1}{N^3} \sum_{n_1, n_2, n_3 \in [N]} \mathbf{P}(T^{n_1} S^{n_2} A \wedge T^{n_3-n_2} S^{n_2} A \wedge T^{n_1} S^{n_3-n_2} A) \\ &= \frac{1}{N^3} \sum_{n_1, n_2, n_3 \in [N]} \mathbf{P}(A \wedge T^{n_3-n_2-n_1} A \wedge S^{n_3-n_2-n_1} A) \\ &\leq \frac{1}{N} \sum_{n=-2N}^N \mathbf{P}(A \wedge T^n A \wedge S^n A). \end{aligned}$$

Thus to prove Theorem B.1 it will suffice to show that

$$\liminf_{N \rightarrow \infty} \mathbf{P} \times \mathbf{P}^{(N)}(A_{12,0,0} \wedge A_{23,0,0} \wedge A_{31,0,0}) > 0.$$

Suppose this were false. Then one can find a sequence $N^{(m)}$ of N , going to infinity as $m \rightarrow \infty$, such that

$$\lim_{m \rightarrow \infty} \mathbf{P} \times \mathbf{P}^{(N^{(m)})}(A_{12,0,0} \wedge A_{23,0,0} \wedge A_{31,0,0}) = 0.$$

By applying Lemma A.15, we can pass to a subsequence if necessary, and obtain a limiting probability measure \mathbf{P}^Δ on $(\Omega^\Delta, \mathcal{B}_{\max}^\Delta)$ with the weak convergence property

$$\lim_{m \rightarrow \infty} \mathbf{P} \times \mathbf{P}^{(N^{(m)})}(E) = \mathbf{P}^\Delta(E) \text{ for all } E \in \mathcal{B}_{\text{reg}}^\Delta.$$

The individual measures $\mathbf{P} \times \mathbf{P}^{(N)}$ can be easily verified to be invariant under triple permutations $(\sigma_1, \sigma_2, \sigma_3)$, and so the limiting measure \mathbf{P}^Δ is also.

By adapting the arguments used to prove Lemma 3.5, one can exploit the above invariance to establish the following relative independence property: If $I_i, I_{i,1}, \dots, I_{i,l}$ are subsets of \mathbf{Z} with $I_i \cap I_{i,1} \cap \dots \cap I_{i,l}$ for $i = 1, 2, 3$, then the factors $\mathcal{B}_{I_1, I_2, I_3}$ and $\bigvee_{j=1}^l \mathcal{B}_{I_{1,j}, I_{2,j}, I_{3,j}}$ are relatively independent conditioning on $\bigvee_{j=1}^l \mathcal{B}_{I_1 \cap I_{1,j}, I_2 \cap I_{2,j}, I_3 \cap I_{3,j}}$, with respect to this probability measure \mathbf{P}^Δ . We omit the details of this as they are essentially the same as in the proof of Lemma 3.5 except for minor notational complications.

We now apply Theorem 4.2 on $(\Omega^\Delta, \mathcal{B}_{\max}^\Delta, \mathbf{P}^\Delta)$, with $J := \{1, 2, 3\}$, $\mathbf{i}_{\max} := \{e : |e| \leq 2\}$, E_e set equal to $A_{ij,0,0}$ if $e = \{i, j\}$ for some $ij = 12, 23, 31$, and $E_e = \Omega^\Delta$ otherwise, and with \mathcal{B}_e set equal to $\mathcal{B}_{I_1, I_2, I_3}$, where I_i is equal to \mathbf{Z} if $i \in e$ and $\mathbf{Z} \setminus \{0\}$ if $i \notin e$. Thus for instance $\mathcal{B}_{\{1,2\}} = \mathcal{B}_{\mathbf{Z}, \mathbf{Z}, \mathbf{Z} \setminus \{0\}}$. The hypotheses of the theorem are easily verified, and by arguing as in the proof of Lemma 4.1 (or Lemma 4.4) we can

find regular events $F'_{i,j} \in \mathcal{B}_{\{i,j\}}$ for $ij = 12, 23, 31$ obeying (11) and

$$\mathbf{P}^\Delta(A_{i,j} \setminus F'_{i,j}) < \eta/10 \text{ for } ij = 12, 23, 31,$$

where $\eta := \mathbf{P}(A) > 0$ is the probability of the original event A . In particular, for all sufficiently large m we have

$$\mathbf{P} \times \mathbf{P}^{N^{(m)}}(A_{i,j} \setminus F'_{i,j}) < \eta/10 \text{ for } ij = 12, 23, 31.$$

Now recall that the variables $n_{1,0}, n_{2,0}, n_{3,0}$ are independently and uniformly distributed on the interval $[N^{(m)}]$. For any fixed $n_{1,0}, n_{2,0}$, the probability that $n_{3,0}$ equals $n_{1,0} + n_{2,0}$ will equal $1/N^{(m)}$ approximately half the time, and 0 the other half of the time. Since the event $A_{1,2} \setminus F'_{1,2}$ is independent of $n_{3,0}$, we thus conclude from Bayes' formula that

$$\mathbf{P} \times \mathbf{P}^{N^{(m)}}(A_{1,2} \setminus F'_{1,2} | n_{3,0} = n_{1,0} + n_{2,0}) < \eta/5$$

for m sufficiently large. Similar arguments in fact give

$$\mathbf{P} \times \mathbf{P}^{N^{(m)}}(A_{i,j} \setminus F'_{i,j} | n_{3,0} = n_{1,0} + n_{2,0}) < \eta/5 \text{ for } ij = 12, 23, 31.$$

On the other hand, from (11) we have¹⁵

$$\mathbf{P} \times \mathbf{P}^{N^{(m)}}(F'_{1,2} \wedge F'_{2,3} \wedge F'_{3,1} | n_{3,0} = n_{1,0} + n_{2,0}) = 0.$$

Combining this with the preceding estimate we see that

$$\mathbf{P} \times \mathbf{P}^{N^{(m)}}(A_{1,2} \wedge A_{2,3} \wedge A_{3,1} | n_{3,0} = n_{1,0} + n_{2,0}) < 3\eta/5.$$

However, the left-hand side equals

$$\mathbf{P}^{N^{(m)}}(\mathbf{P}(T^{n_{1,0}} S^{n_{2,0}} A \wedge T^{n_{3,0}-n_{2,0}} S^{n_{2,0}} A \wedge T^{n_{1,0}} S^{n_{3,0}-n_{1,0}} A) | n_{3,0} = n_{1,0} + n_{2,0})$$

which simplifies (using the shift invariance) to $\mathbf{P}(A) = \eta$. Thus we have $\eta < 3\eta/5$, a contradiction. This proves Theorem B.1. \blacksquare

Remark B.4. At present, the hypergraph regularity method is known to yield the Furstenberg-Katznelson recurrence theorem, but more powerful recurrence theorems, such as the Bergelson-Leibman polynomial recurrence theorem, the Furstenberg-Katznelson IP-Szemerédi theorem, and the Furstenberg-Katznelson density Hales-Jewett theorem, have not yet been successfully obtained by this method (either in the finitary or infinitary settings). It is not clear to the author whether this represents any fundamental limitations to the method. A possible test problem would be the refinement of Szemerédi's theorem that the set of possible differences amongst the arithmetic progressions of a given length is syndetic (has bounded gaps); this was established for instance in [7] by ergodic methods but does not currently have a non-ergodic proof.

¹⁵Note how important it is here that the event $F'_{1,2} \wedge F'_{2,3} \wedge F'_{3,1}$ be *empty*, rather than merely being a null event with respect to \mathbf{P}^Δ . In the latter case, the event would have a small but nonzero measure in $\mathbf{P} \times \mathbf{P}^{N^{(m)}}$, and we would be unable to condition this event to the vanishingly small probability event $n_{3,0} = n_{1,0} + n_{2,0}$ without losing control on the conditional probability. The point is that the constraint $n_{3,0} = n_{1,0} + n_{2,0}$ creates a "diagonal measure" which is singular with respect to \mathbf{P}^Δ , and so null events in \mathbf{P}^Δ do not necessarily restrict to null events on the diagonal measure. However, events which have empty intersection with respect to \mathbf{P}^Δ will clearly continue to have empty intersection with respect to the diagonal measure. This robustness with respect to change of measure is what makes Theorem 4.2 (which is basically a mechanism for converting null events to empty events) so powerful.

REFERENCES

- [1] M. Ajtai, E. Szemerédi, *Sets of lattice points that form no squares*, Studia Scientiarum Mathematicarum Hungarica. **9** (1974), 9–11.
- [2] N. Alon, A. Shapira, *A characterization of the (natural) graph properties testable with one-sided error*, preprint.
- [3] F. Chung, R. Graham, *Quasi-random hypergraphs*, Random Struct Alg **1**(1) (1990), 105–124.
- [4] P. Erdős, P. Frankl, V. Rödl, *The asymptotic number of graphs not containing a fixed subgraph and a problem for hypergraphs having no exponent*, Graphs Combin. **2** (1986), no. 2, 113–121.
- [5] P. Frankl, V. Rödl, *The uniformity lemma for hypergraphs*, Graphs Combinat. **8**(4) (1992), 309–312.
- [6] P. Frankl, V. Rödl, *Extremal problems on set systems*, Random Struct. Algorithms **20** (2002), no. 2, 131–164.
- [7] H. Furstenberg, *Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions*, J. Analyse Math. **31** (1977), 204–256.
- [8] H. Furstenberg, *Recurrence in Ergodic theory and Combinatorial Number Theory*, Princeton University Press, Princeton NJ 1981.
- [9] H. Furstenberg, Y. Katznelson, *An ergodic Szemerédi theorem for commuting transformations*, J. Analyse Math. **34** (1978), 275–291.
- [10] H. Furstenberg, Y. Katznelson, D. Ornstein, *The ergodic theoretical proof of Szemerédi’s theorem*, Bull. Amer. Math. Soc. **7** (1982), 527–552.
- [11] J.-Y. Girard, *Herbrand’s theorem and proof theory*, Proceedings of the Herbrand symposium (Marseilles, 1981), 29–38, Stud. Logic Found. Math., 107, North-Holland, Amsterdam, 1982.
- [12] T. Gowers, *Hypergraph regularity and the multidimensional Szemerédi theorem*, preprint.
- [13] B. Green, T. Tao, *The primes contain arbitrarily long arithmetic progressions*, preprint.
- [14] Host, B., *Progressions arithmétiques dans les nombres premiers (d’après B. Green and T. Tao)*, Séminaire Bourbaki, Mars 2005, 57eme année, 2004-2005, no. 944.
- [15] B. Host, B. Kra, *Nonconventional ergodic averages and nilmanifolds*, Ann. Math. **161** (2005), 397–488.
- [16] J. Komlós, M. Simonovits, *Szemerédi’s regularity lemma and its applications in graph theory*, Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993), 295–352, Bolyai Soc. Math. Stud., 2, János Bolyai Math. Soc., Budapest, 1996.
- [17] Kra, B., *The Green-Tao Theorem on arithmetic progressions in the primes: an ergodic point of view*, to appear, Bulletin of Amer. Math. Soc.
- [18] L. Lovász, B. Szegedy, *Limits of dense graph sequences*, preprint.
- [19] B. Nagle, V. Rödl, M. Schacht, *The counting lemma for regular k -uniform hypergraphs*, to appear, Random Structures and Algorithms.
- [20] V. Rödl, M. Schacht, *Regular partitions of hypergraphs*, preprint.
- [21] V. Rödl, M. Schacht, E. Tengan, N. Tokushige, *Density theorems and extremal hypergraph problems*, preprint.
- [22] V. Rödl, J. Skokan, *Regularity lemma for k -uniform hypergraphs*, to appear, Random Structures and Algorithms.
- [23] V. Rödl, J. Skokan, *Applications of the regularity lemma for uniform hypergraphs*, preprint.
- [24] I. Ruzsa, E. Szemerédi, *Triple systems with no six points carrying three triangles*, Colloq. Math. Soc. J. Bolyai **18** (1978), 939–945.
- [25] J. Solymosi, *Note on a generalization of Roth’s theorem*, Discrete and computational geometry, 825–827, Algorithms Combin. **25**, Springer Verlag, 2003.
- [26] E. Szemerédi, *On sets of integers containing no four elements in arithmetic progression*, Acta Math. Acad. Sci. Hungar. **20** (1969), 89–104.
- [27] E. Szemerédi, *Regular partitions of graphs*, in “Problèmes Combinatoires et Théorie des Graphes, Proc. Colloque Inter. CNRS,” (Bermond, Fournier, Las Vergnas, Sotteau, eds.), CNRS Paris, 1978, 399–401.
- [28] T. Tao, *Szemerédi’s regularity lemma revisited*, preprint.
- [29] T. Tao, *A quantitative ergodic theory proof of Szemerédi’s theorem*, preprint.
- [30] T. Tao, *A variant of the hypergraph removal lemma*, preprint.
- [31] T. Tao, *The Gaussian primes contain arbitrarily shaped constellations*, preprint.
- [32] T. Tao, *An ergodic transference theorem*, unpublished.

- [33] P. Varnavides, *On certain sets of positive density*, J. London Math. Soc. **39** (1959), 358–360.

DEPARTMENT OF MATHEMATICS, UCLA, LOS ANGELES CA 90095-1555

E-mail address: `tao@math.ucla.edu`